

Subgroup decomposition in $\text{Out}(F_n)$

Part I: Geometric Models

Michael Handel and Lee Mosher

February 14, 2013

Abstract

This is the first in a series of four papers, announced in [HM13a], that develop a decomposition theory for subgroups of $\text{Out}(F_n)$.

In this paper we develop further the theory of geometric EG strata of relative train track maps originally introduced in [BFH00] Section 5, with our focus trained on certain 2-dimensional models of such strata called “geometric models” and on the interesting properties of these models. A secondary purpose of this paper is to serve as a central reference for the whole series regarding basic facts of the general theory for elements of $\text{Out}(F_n)$.

An outer automorphism $\phi \in \text{Out}(F_n)$ is said to be *geometric* if there exists a compact surface S with nonempty boundary ∂S , a homeomorphism $\Phi: S \rightarrow S$, and an isomorphism between $\pi_1(S)$ and F_n , such that the outer automorphism of F_n induced by the homeomorphism Φ equals ϕ . Any such surface S deformation retracts to a finite graph G called a *spine*. The homeomorphism Φ of S induces a homotopy equivalence h of G , and either of Φ or h may be considered as a topological representative of the outer automorphism ϕ . If ∂S is connected and the mapping class of Φ is pseudo-Anosov then ϕ is *fully irreducible*, meaning that for each proper, nontrivial free factor $A < F_n$, the conjugacy class of A in F_n is not ϕ -periodic. Conversely, as proved in [BH92] using train track maps, if ϕ is fully irreducible and if ϕ has a periodic conjugacy class then ϕ is geometric.

The concept of geometricity also arises in the context of strata of relative train track maps, as follows. Consider a compact surface S with nonempty boundary ∂S and a homeomorphism $\Psi: S \rightarrow S$ with pseudo-Anosov mapping class. We make no assumption that ∂S is connected, although we single out one component of ∂S as the “top boundary”, the remaining ones being called the “lower boundary components”. Consider also a graph G' , a homotopy equivalence $f': G' \rightarrow G'$, and a homotopically nontrivial gluing map from each lower boundary component to G' such that the gluing maps are equivariant with respect to the action of Ψ on the lower components of ∂S and the action of h of H . The result of this gluing is a 2-complex Y and a self homotopy equivalence $h: Y \rightarrow Y$. Any such 2-complex deformation retracts to a spine G' that includes G'' as a subcomplex, and the homotopy equivalence $h: Y \rightarrow Y$ induces a homotopy equivalence $f'': G'' \rightarrow G''$ which extends f' . Following [BFH00] Definition 5.1.4, in Definition 2.2 below we use this construction to define geometricity of an EG stratum of a relative train track map $f: G \rightarrow G$ representing some $\phi \in \text{Out}(F_n)$: the subgraphs $G' = G_{r-1} \subset G_r = G''$ will be adjacent filtration elements, the

maps f', f'' are the restrictions of f , and the stratum $H_r = G_r \setminus G_{r-1}$ is, by definition, a geometric EG stratum of f . The 2-complex Y and the homotopy equivalence $h: Y \rightarrow Y$ define what we shall call a *weak geometric model for f with respect to the stratum H_r* , the adjective “weak” emphasizing that it is a topological model only for the restricted homotopy equivalence $f|G_r$. A (full) *geometric model* for f with respect to H_r , which is indeed a topological model for all of f , is then defined by attaching the higher strata $G \setminus G_r$ to Y to form a 2-complex X homotopy equivalent to G , and extending $h: Y \rightarrow Y$ appropriately to a homotopy equivalence of X .

Geometricity of EG strata plays an explicit role in the proof of the Tits alternative for $\text{Out}(F_n)$ in [BFH00], with weak geometric models often playing an implicit supporting role. In this paper we give weak geometric models and (full) geometric models a central role, developing their theory in some depth, and obtaining applications needed for later in this series.

Geometricity as an invariant of laminations. The decomposition theory for $\text{Out}(F_n)$ developed on [BFH00] associates to each $\phi \in \text{Out}(F_n)$ a finite set $\mathcal{L}(\phi)$ of “attracting laminations” (see Sections 1.3 and 1.6 for a review). After passing to a power of ϕ which is “rotationless”, and using the very nicest kind of topological representative $f: G \rightarrow G$ called a “completely split relative train track map” or just a CT, there is a natural correspondence between the set $\mathcal{L}(\phi)$ and the set of EG strata of f , where $\Lambda \in \mathcal{L}(\phi)$ corresponds to H_r if and only if iterating the edges of H_r produces the leaves of Λ in the limit. The elements of the $\mathcal{L}(\phi)$ can be distinguished by their “free factor supports”, and the elements of the two sets $\mathcal{L}(\phi)$ and $\mathcal{L}(\phi^{-1})$ can be put in bijective correspondence where $\Lambda^+ \in \mathcal{L}(\phi)$ corresponds to $\Lambda^- \in \mathcal{L}(\phi^{-1})$ if and only if Λ^+, Λ^- have the same free factor supports, in which case we say that this pair forms a *dual lamination pair*.

Although geometricity is defined in [BFH00] and here in Definition 2.2 as an invariant of EG strata of relative train tracks, the following result says that it is actually an invariant of laminations, and furthermore it is an invariant of laminations pairs.

Proposition. (*Proposition 2.13*) *For any $\phi \in \text{Out}(F_n)$ and $\Lambda^+ \in \mathcal{L}(\phi)$, if Λ^+ corresponds to a geometric EG stratum in some CT representing some rotationless power of ϕ , then Λ^+ corresponds to a geometric EG stratum in every CT representing every rotationless power of ϕ . Furthermore, if this holds, and if $\Lambda^- \in \mathcal{L}(\phi^{-1})$ is dual to Λ^+ , then Λ^- corresponds to a geometric EG stratum of every CT representing every rotationless power of ϕ^{-1} .*

The full statement of this proposition gives an explicit property, expressed solely in terms of ϕ and Λ without regard to a representative CT, which nonetheless characterizes geometricity of the EG stratum associated to Λ in any CT representing any rotationless power of ϕ . This property says that there is a finite ϕ -invariant set of conjugacy classes in F_n whose free factor support equals the free factor support of Λ .

Complementarity properties of geometric models. Consider a closed surface S decomposed as a union of two compact connected subsurfaces $S = S_1 \cup S_2$ such that $S_1 \cap S_2 = \partial S_1 = \partial S_2 = c$ is connected and neither of S_1, S_2 is a disc. In this setting the subgroup of $\mathcal{MCG}(S)$ that stabilizes S_1 (up to isotopy) is the same as the subgroup that

stabilizes S_2 and the subgroup that stabilizes \mathcal{C} preserving transverse orientation. Furthermore, one can understand this subgroup rather explicitly: it is the quotient of direct sum of the relative mapping class groups $\mathcal{MCG}(S_1, \partial S_1)$ and $\mathcal{MCG}(S_2, \partial S_2)$, modulo identification of Dehn twists supported in collar neighborhoods of c in S_1 and in S_2 .

By contrast, such behavior fails in the setting of $\text{Out}(F_n)$ when using a free factorization $F_n = A * B$ in place of the decomposition $S = S_1 \cup S_2$: the subgroup of $\text{Out}(F_n)$ that preserves the conjugacy class of the free factor A is not equal to the subgroup that preserves the conjugacy class of B . One way to think of this is that there is no well-defined way to determine a complement of a free factor of F_n amongst other free factors (up to conjugacy).

We study a version of complementarity which works in the context of geometric models. For purposes of brevity, in this introduction we describe this phenomenon only in the special case that H_r is the top stratum of a CT $f: G \rightarrow G$ representing a rotationless $\phi \in \text{Out}(F_n)$, in which case the full geometric model X and the weak geometric model Y are the same. In this special case, under the quotient map $G_{r-1} \cup S \mapsto X$, the interior of the surface S embeds as an open subset of X whose complement is a finite graph L which is the disjoint union of G_{r-1} and a circle which is as a copy of the top boundary circle $\partial_0 S$; we refer to the graph L as the *complementary graph* of the geometric model. One may reconstruct X as the quotient obtained from the disjoint union $L \cup S$ by gluing the lower boundary circles to G_{r-1} as before and identifying the upper boundary circle $\partial_0 S$ in S with its copy in L .

Like subgraphs of marked graphs, the inclusion of L into X is π_1 -injective on each non-contractible component, and the image subgroups are each malnormal and are “mutually” malnormal (see Lemma 2.7). The collection of conjugacy classes of these image subgroups forms what we call a “subgroup system” in $F_n \approx \pi_1(X)$, denoted $[\pi_1 L]$ (see Section 1.1.2 for “subgroup systems”). Also, the restriction to S of the quotient map $L \cup S \mapsto X$ is π_1 -injective and the conjugacy class of the image of $\pi_1(S)$ is a subgroup system in F_n denoted $[\pi_1 S]$. In general the subgroup systems $[\pi_1 L]$ and $[\pi_1 S]$ do not behave as nicely as free factor systems, and in fact $[\pi_1 L]$ is *never* a free factor system in F_n , although $[\pi_1 S]$ may or may not be a free factor system depending on the example.

The subgroup systems $[\pi_1 L]$ and $[\pi_1 S]$, and the subgroups of $\text{Out}(F_n)$ that stabilize them, will play an important role in later applications in this series of papers. For example, in the present context where we have assumed H_r to be the top stratum, the conjugacy classes of F_n represented by circuits in L are precisely the conjugacy classes that are not weakly attracted to the lamination Λ^+ under iteration of ϕ ; this fact is a key observation underlying the Weak Attraction Theorem of [BFH00], and its generalization dropping the assumption on H_r will play a role in our general study of weak attraction theory in Part III [HM13c]. Also, in Part II [HM13b] we shall need to study nonrotationless elements of $\text{Out}(F_n)$ which stabilize $[\pi_1 L]$, such as certain nonrotationless roots of ϕ . Also, regarding $[\pi_1 S]$, an important step in the proof of Proposition 2.13 stated above is to prove that the free factor support of the lamination Λ^+ equals the free factor support of $[\pi_1 S]$.

The following result, which is perhaps the most significant new result of this paper, says roughly that if an outer automorphism preserves the complementary subgraph L then it preserves the surface S . To state it, we use that the natural action of $\text{Aut}(F_n)$ on subgroups induces natural actions of $\text{Out}(F_n)$ on conjugacy classes of subgroups and on subgroup systems, with respect to which we may define the subgroup $\text{Stab}(\mathcal{A}) < \text{Out}(F_n)$ which stabilizes a subgroup system \mathcal{A} .

Proposition (Proposition 2.14, special case). *For any $\theta \in \text{Out}(F_n)$, if $\theta[\pi_1 L] = [\pi_1 L]$ then $\theta[\pi_1 S] = [\pi_1 S]$, inducing a well-defined outer automorphism $\theta \mid \pi_1 S \in \text{Out}(\pi_1 S)$. Furthermore, this outer automorphism preserves the set of conjugacy classes associated to the components of ∂S and so, by the Dehn-Nielsen-Baer theorem, $\theta \mid \pi_1(S)$ lives in the natural $\mathcal{MCG}(S)$ subgroup of $\text{Out}(\pi_1 S)$. One obtains by this process a homomorphism $\text{Stab}[\pi_1 L] \mapsto \mathcal{MCG}(S)$.*

This proposition is based on Lemma 2.15 we can be regarded as a topological version of the proposition, stated in terms of self homotopy equivalences of a geometric model. In Lemma 2.18 we prove an important addendum to Lemma 2.15, which will be used in Part II [HM13b] when we analyze elements of $\text{Out}(F_n)$ that act trivially on homology with $\mathbb{Z}/3$ coefficients. In the proof of Proposition 2.14, the application of the Dehn-Nielsen-Baer theorem comes after showing that $\theta \mid \pi_1 S$ preserves the conjugacy classes of $\pi_1 S$ associated to the boundary circles of S . Although the gluing of lower boundary circles of S to G_{r-1} will in general not be 1–1, it may happen that certain components $c \subset \partial S$ remain as “free boundary circles” in the geometric model, by which we mean that some half-open annulus neighborhood of c in S embeds with image an open subset of the geometric model. For example in the case where H_r is a top stratum: the top boundary circle $\partial_0 S$ is always free; and a lower boundary circle is free if it is identified homeomorphically to a component of G_{r-1} and has no point identified with any point on another lower boundary circle. Lemma 2.18 says that $\theta \mid \pi_1(S) \in \mathcal{MCG}(S)$ preserves the conjugacy classes of $\pi_1 S$ associated to the subset of free boundary circles.

Vertex group systems. As mentioned above, the subgroup system $[\pi_1 L]$ associated to complementary subgroup L of a geometric model X is never be a free factor system, as we show in Lemma 2.5 (4). In the special case of a geometric outer automorphism, where $X = S$ is a compact surface and $L = \partial S \neq \emptyset$, this follows from a result due to Stallings [Sta00] which shows that for any free factorization $\pi_1(S) = A_1 * \dots * A_K$ there exists a component of ∂S which is not conjugate into any factor A_k , and our argument is an adaptation of Stallings’ proof.

Subgroup systems in F_n can be ill-behaved in general. We shall show that the subgroup system $[\pi_1 L]$ has an important special property. Given a minimal action of F_n on an \mathbf{R} -tree T such that the stabilizer of each arc is trivial, there are only finitely many F_n -orbits of points $x \in T$ such that the stabilizer of x is nontrivial [GL95]. The collection of nontrivial point stabilizers of T is closed under conjugacy, and the corresponding subgroup system is called the *vertex group system* of T . Vertex group systems, while not quite as well-behaved as free factor systems, nonetheless satisfy some useful properties; in Proposition 3.2 we give a proof that we learned from Mark Feighn of a descending chain condition on vertex group systems.

Our main result relating geometric models and vertex group systems is a more general version of the following proposition, and it will be used for various purposes in later parts of this series.

Proposition. (Proposition 3.3, special case) *For any geometric model X , the subgroup system $[\pi_1 L]$ is the vertex group system of some minimal action of F_n on an \mathbf{R} -tree with trivial arc stabilizers.*

Description of the contents. Section 1 is a detailed but terse review of the decomposition theory for individual elements of $\text{Out}(F_n)$ as developed in [BFH00] and in [FH11]. This section is designed as a central reference, and a reader who is conversant in the language of relative train track and outer automorphisms as presented in [BFH00] and [FH11] may only need to refer to Section 1 occasionally.

Section 2 contains the results about geometricity of EG strata and of attracting laminations and about the geometric models. The dichotomy between geometric and nongeometric EG strata is used in the statements and proofs of many results in [BFH00], and the same is true in this series; see for example in Theorem C of the introductory paper of the series [HM13a].

Section 3 is a study of vertex group systems and their relation to geometric strata.

Contents

1 Preliminaries: Decomposition of outer automorphisms of free groups	6
1.1 F_n and its subgroups, marked graphs, and lines.	6
1.1.1 The geometry of F_n and its subgroups.	6
1.1.2 Subgroup systems and free factor systems.	8
1.1.3 Restrictions of outer automorphisms.	9
1.1.4 Marked graphs, paths, and circuits.	9
1.1.5 Lines and rays.	11
1.1.6 The path maps $f_\#$ and $f_{\#}$, and bounded cancellation.	12
1.2 Subgroup systems carrying lines and other things.	14
1.2.1 Subgroup systems carrying lines and rays.	14
1.2.2 Free factor supports of lines, rays, and subgroup systems.	15
1.3 Attracting laminations.	17
1.4 Principal automorphisms and rotationless outer automorphisms.	18
1.5 Relative train track maps and CTs	21
1.5.1 Definitions of relative train track maps and CTs	21
1.5.2 Facts about CTs, their Nielsen paths, and their zero strata.	25
1.5.3 Facts about principal lifts, principal directions, and principal rays. .	29
1.5.4 Properties of EG strata.	31
1.6 Properties of Attracting Laminations	33
1.6.1 The relation between EG strata and attracting laminations.	34
1.6.2 Tiles and applications.	34
1.6.3 EG principal rays and Fix_N	36
1.6.4 Pushing forward attracting laminations.	38
2 Geometric EG strata and geometric laminations	41
2.1 Defining and characterizing geometric strata	42
2.1.1 Defining weak geometric models and geometric strata.	42
2.1.2 Defining geometric models.	44
2.1.3 Invariant free factor systems associated to a geometric model.	46
2.1.4 Characterizing geometric strata: Proof of Fact 2.3	48
2.2 Complementary subgraph and peripheral splitting	51

2.3	The laminations of a geometric stratum	54
2.3.1	Review of Nielsen–Thurston theory.	54
2.3.2	Comparing free group laminations and surface laminations.	56
2.3.3	Application: The inclusion lattice of attracting laminations.	58
2.4	Geometricity is an invariant of a dual lamination pair	59
2.5	Stabilizing its complement also stabilizes the surface	60
2.6	Preserving the free boundary circles.	63
3	Vertex groups and vertex group systems	66
3.1	Vertex group systems	67
3.2	Geometric models and vertex group systems	68

1 Preliminaries: Decomposition of outer automorphisms of free groups

In this section we give a self-contained account of basic results about $\text{Out}(F_n)$ needed throughout this series of papers, citing outside sources, particularly [BFH00] and [FH11], for the reader to verify the correctness of statements. With that intent, we try to be as succinct as possible. We shall state numerous “Facts” which are either citations from the literature, or are quickly proved using results from the literature. We state also some “Lemmas” whose proofs require a bit more care.

A reader familiar with the results of [BFH00] and [FH11] may find it convenient to use this section more as a central reference for this whole series of papers, rather than as a prerequisite for reading the series.

1.1 F_n and its subgroups, marked graphs, and lines.

For $n \geq 2$ let R_n denote the rose with n directed edges labelled E_1, \dots, E_n . Let F_n denote the rank n free group $\pi_1(R_n)$, with E_1, \dots, E_n as a free basis, and fix a lifted base point in the universal cover \tilde{R}_n thereby identifying F_n as the group of deck transformations, and so \tilde{R}_n is the Cayley graph of F_n with respect to that free basis.

Let $\text{Aut}(F_n)$ be the automorphism group of F_n , let $\text{Inn}(F_n)$ be the group of inner automorphisms of F_n , and let $\text{Out}(F_n) = \text{Aut}(F_n)/\text{Inn}(F_n)$ be the outer automorphism group. The action of $\text{Aut}(F_n)$ on elements and subgroups of F_n induces actions of $\text{Out}(F_n)$ on conjugacy classes of elements and of subgroups.

We use the notation $[\cdot]$ to denote conjugacy classes in F_n of elements and of subgroups, i.e. F_n -orbits under the action of F_n on itself by inner automorphisms. We use $\mathcal{C} = \mathcal{C}(F_n)$ to denote the set of conjugacy classes of nontrivial elements of F_n .

1.1.1 The geometry of F_n and its subgroups.

Recall that in a finitely generated group G , a finitely generated subgroup $H < G$ is *undistorted* if the inclusion map $H \hookrightarrow G$ is a quasi-isometric embedding with respect to the word metrics d_H, d_G .

Fact 1.1. *For any finite rank subgroup $A < F$ the following hold:*

- (1) *A has finite index in its own normalizer.*
- (2) *A is undistorted in F .*

Proof. Under the action $F_n \curvearrowright \tilde{R}_n$, the restricted action of A on its minimal subtree T is cocompact. It follows that A has finite index in the subgroup of F_n that stabilizes T , but that subgroup is the normalizer of A , proving (1). The injection $T \hookrightarrow \tilde{R}_n$ is an isometric embedding, and for any $x \in T$ the orbit map $A \rightarrow T$ defined by $a \mapsto a \cdot x$ is a quasi-isometric embedding, proving (2). \square

This fact has the following consequences which we use henceforth without comment, regarding the natural identification between the Gromov boundary ∂F_n and the Cantor set of ends of the tree \tilde{R}_n . For any finite rank subgroup $A < F_n$, it follows by item (2) that the Gromov boundary ∂A is naturally identified with the closure in ∂F_n of the set of fixed point pairs for the action of nontrivial elements of A ; in particular when $A = \langle \gamma \rangle$ is infinite cyclic, ∂A is a 2-point subset of ∂F_n , namely the fixed point set for the action of γ . Also, under the action of F_n on \tilde{R}_n , or under any properly discontinuous cocompact action of F_n on a simplicial tree T whose boundary ∂T has been identified with ∂F_n , it follows by item (1) that the action of A on the unique minimal A -invariant subtree T^A of T is cocompact; the Gromov boundary of T^A is identified with ∂A , and T^A equals the *convex hull* of ∂A denoted $\mathcal{H}(\partial A)$, which is just the union of all lines in T whose endpoints are in ∂A .

The following is a standard fact about undistorted torsion free subgroups of word hyperbolic groups, but in our context it has a rather low technology proof.

Fact 1.2. *For any finitely generated subgroups $A_1, A_2 < F_n$, their intersection $A_1 \cap A_2$ is the trivial subgroup if and only if $\partial A_1 \cap \partial A_2 = \emptyset$. More generally, $\partial(A_1 \cap A_2) = \partial A_1 \cap \partial A_2$.*

Proof. Let $p \in \tilde{R}_n$ be the base vertex of the Cayley graph. Let $|\cdot|$ denote the word norm in F_n and let $d(\cdot, \cdot)$ denote the word metric, so $d(g, h) = d(g(p), h(p)) = |h^{-1}g|$. The inclusion $\partial(A_1 \cap A_2) \subset \partial A_1 \cap \partial A_2$ being obvious, consider $\xi \in \partial A_1 \cap \partial A_2$. In \tilde{R}_n the intersection $\mathcal{H}(\partial A_1) \cap \mathcal{H}(\partial A_2)$ contains a ray $[x, \xi)$. Choose a base vertex $p_i \in \mathcal{H}(\partial A_i)$, $i = 1, 2$ and $g_i \in F_n$ so that $g_i(p) = p_i$. By cocompactness of the action of A_i on $\mathcal{H}(\partial A_i)$, $i = 1, 2$, there exists $C \geq 0$, vertex sequences $x_k \in \mathcal{H}(\partial A_1)$ and $y_k \in \mathcal{H}(\partial A_2)$, and sequences $h_{1k} \in A_1$, $h_{2k} \in A_2$, such that $h_{1k}(p_1) = x_k$, $h_{2k}(p_2) = y_k$, $d(x_k, y_k) \leq C$ in \tilde{R}_n , $d(x_k, [x, \xi)), d(y_k, [x, \xi)) \leq C$, and as $k \rightarrow +\infty$ both of the sequences x_k, y_k approach ξ in the compactified Cayley graph $\tilde{R}_n \amalg \partial F_n$. It follows that

$$\begin{aligned} |h_{1k}^{-1}h_{2k}| &= d(h_{1k}(p), h_{2k}(p)) \\ &\leq d(h_{1k}(p), h_{1k}(p_1)) + d(h_{1k}(p_1), h_{2k}(p_2)) + d(h_{2k}(p_2), h_{2k}(p)) \\ &= d(p, p_1) + d(x_k, y_k) + d(p_2, p) \end{aligned}$$

is bounded independent of k . Since F_n has only finitely many elements with a given bound on word length, passing to a subsequence we may assume that $h_{1k}^{-1}h_{2k}$ is constant. In particular $h_{1,0}^{-1}h_{2,0} = h_{1k}^{-1}h_{2k}$ and so $h_{1k}h_{1,0}^{-1} = h_{2k}h_{2,0}^{-1} = f_k \in A_1 \cap A_2$ for all k . The distance between $x_k = h_{1k}(p_1)$ and $f_k(p_1) = h_{1k}(h_{1,0}(p_1))$ is uniformly bounded for all k ,

and so $f_k(p_1)$ converges to ξ in $\tilde{R}_n \cup \partial F_n$. Truncating an initial sequence we may assume f_k is nontrivial for all k . Letting $\eta_k^+ \in \partial(A_1 \cap A_2)$ be the attracting endpoint of the axis of f_k , it follows that $\eta_k^+ \rightarrow \xi$ as $k \rightarrow +\infty$. Since $\partial(A_1 \cap A_2)$ is closed it follows that $\xi \in \partial(A_1 \cap A_2)$. \square

1.1.2 Subgroup systems and free factor systems.

We recall from [BFH00] Section 2.6 the concepts of free factor systems. We also introduce a generalized notion called “subgroup systems”, to lay the ground for the “vertex group systems” introduced in Section 3 and for the “fixed subgroup systems” that arise in the study of rotationless outer automorphisms and that are used in Part II [HM13b].

Define a *subgroup system* in F_n to be a finite set \mathcal{K} each of whose elements is a conjugacy class of finite rank subgroups of F_n . The individual elements of the set \mathcal{K} are called its *components*. Note that a subgroup system \mathcal{K} is completely determined by the set of subgroups $K < F_n$ such that $[K] \in \mathcal{K}$. With this in mind, we often abuse notation by writing $K \in \mathcal{K}$ when we really mean $[K] \in \mathcal{K}$. There is a partial order on subgroup systems denoted $\mathcal{K}_1 \sqsubset \mathcal{K}_2$ and defined by requiring that for each subgroup $K \in \mathcal{K}_1$ there exists a subgroup $K' \in \mathcal{K}_2$ such that $K < K'$. We refer to the relation $\mathcal{K}_1 \sqsubset \mathcal{K}_2$ by saying for example that “ \mathcal{K}_1 is contained in \mathcal{K}_2 ” or that “ \mathcal{K}_2 carries \mathcal{K}_1 ”. In particular we say that \mathcal{K}_1 is *properly contained* in \mathcal{K}_2 if $\mathcal{K}_1 \sqsubset \mathcal{K}_2$ and $\mathcal{K}_1 \neq \mathcal{K}_2$. The action of $\text{Out}(F_n)$ on the set of conjugacy classes of subgroups induces an action on the set of subgroup systems. The *stabilizer* of a subgroup system \mathcal{K} is the subgroup of all $\theta \in \text{Out}(F_n)$ such that $\theta(\mathcal{K}) = \mathcal{K}$.

A *free factor system* in F_n is a (possibly empty) subgroup system of the form $\mathcal{F} = \{[A_1], \dots, [A_k]\}$ such that there exists a free factorization $F_n = A_1 * \dots * A_k * B$ where each A_i is nontrivial; the subgroup B may or may not be trivial. The partial order \sqsubset defined generally on subgroup systems is consistent with the partial order on free factor systems given in the introduction. The action of $\text{Out}(F_n)$ on the set of subgroup systems restricts to an action on the set of free factor systems.

We often take the liberty of referring to a free factor system $\{[A]\}$ with a single component as a *free factor*, rather than the lengthier but more appropriate locution “free factor conjugacy class”. This should not cause any confusion since each element of a free factor system is a conjugacy class of subgroups of F_n as opposed to a subgroup itself.

The following fact and definition taken from Section 2 of [BFH00] is essentially a refinement of the Kurosh Subgroup Theorem:

Fact 1.3 (Definition of the meet of free factor systems). *Every collection $\{\mathcal{F}_i\}_{i \in I}$ of free factor systems has a well-defined meet $\wedge\{\mathcal{F}_i\}$, which is the unique maximal free factor system \mathcal{F} such that $\mathcal{F} \sqsubset \mathcal{F}_i$ for each $i \in I$. Furthermore, for any free factor $A < F_n$ we have $[A] \in \wedge\{\mathcal{F}_i\}$ if and only if there exists an indexed collection of subgroups $\{A_i\}_{i \in I}$ such that $[A_i] \in \mathcal{F}_i$ for each i and $A = \cap_{i \in I} A_i$. In particular it follows that the relation \sqsubset satisfies the descending chain condition.*

Malnormal subgroup systems. Recall that a subgroup $K < F_n$ is *malnormal* if $K^x \cap K$ is trivial for all $x \in F_n - K$; note in particular that a malnormal subgroup is its own normalizer.

More generally, define a subgroup system \mathcal{K} to be *malnormal* if each of its subgroups $K \in \mathcal{K}$ is malnormal and for all $K, K' \in \mathcal{K}$, if $K \cap K'$ is nontrivial then $K = K'$. Every free factor $A \in F_n$ is malnormal, and every free factor system \mathcal{F} is malnormal; this can be recovered by applying Fact 1.3 to the meet of the collection $\{\mathcal{A}\}$. By Fact 1.2, if a subgroup system \mathcal{K} is malnormal then for any two subgroups $K, K' \in \mathcal{K}$, if $K \neq K'$ then $\partial K \cap \partial K' = \emptyset$.

1.1.3 Restrictions of outer automorphisms.

Consider $\theta \in \text{Out}(F_n)$ and a finite rank subgroup $A < F_n$ whose conjugacy class $[A]$ is fixed by the action of θ . We may choose $\Theta \in \text{Aut}(F_n)$ representing θ so that $\Theta(A) = A$, and we obtain $\Theta \mid A \in \text{Aut}(A)$. If the outer automorphism class of $\Theta \mid A$ is well-defined in the group $\text{Out}(A)$ independent of the choice of Θ , then we denote this by $\theta \mid A \in \text{Out}(A)$ and we say that $\theta \mid A$ is *well-defined*. This need not always happen, but the following gives a sufficient condition:

Fact 1.4. *Given a finite rank subgroup $A < F_n$ which is its own normalizer in F_n , for any $\theta \in \text{Out}(F_n)$ such that $\theta[A] = [A]$, the restriction $\theta \mid A \in \text{Out}(A)$ is well-defined, and the map $\theta \mapsto \theta \mid A$ defines a homomorphism $\text{Stab}[A] \hookrightarrow \text{Out}(A)$. If moreover A is malnormal then the conjugacy class in A of an element or subgroup of A is fixed by $\theta \mid A$ if and only if the conjugacy class in F_n determined by that element or subgroup is fixed by θ .*

Proof. Let $\Theta' \in \text{Aut}(F_n)$ be another representative of θ such that $\Theta'(A) = A$. Choose $d \in F_n$ with associated inner automorphism $\iota_d(\gamma) = d\gamma d^{-1}$ so that $\Theta' = \iota_d \circ \Theta$. It follows that $A = \Theta'(A) = \iota_d(\Theta(A)) = \iota_d(A)$, and so d normalizes A , implying that $d \in A$. The equation $\Theta' = \iota_d \circ \Theta$ therefore implies that Θ, Θ' represent the same outer automorphism of A . Given $\theta_1, \theta_2 \in \text{Stab}[A]$, and choosing representatives $\Theta_1, \Theta_2 \in \text{Aut}(F_n)$ each preserving the subgroup A , it follows that $\Theta_1 \cdot \Theta_2 \in \text{Aut}(F_n)$ also preserves A ; since $\Theta_1 \cdot \Theta_2$ represents the element $\theta_1 \cdot \theta_2$ in $\text{Stab}[A]$, and since it also represents the element $(\theta_1 \mid A) \cdot (\theta_2 \mid A)$ in $\text{Out}(A)$, it follows that the map $\text{Stab}[A] \hookrightarrow \text{Out}(A)$ takes $\theta_1 \cdot \theta_2$ to $(\theta_1 \mid A) \cdot (\theta_2 \mid A)$ and is therefore a homomorphism.

Malnormality of A implies that the injection $A \hookrightarrow F_n$ induces injections from the sets of conjugacy classes in A of elements and subgroups to the sets of conjugacy classes in F_n of elements and subgroups, and these injections are natural with respect to the actions of $\theta \mid A$ on conjugacy classes in A and of θ on conjugacy classes in F_n . The final “If moreover...” sentence follows. \square

1.1.4 Marked graphs, paths, and circuits.

The subset theoretic difference operator, for any set A and subset B , is denoted throughout as $A - B = \{x \in A \mid x \notin B\}$. We also need a subgraph difference operator, and more generally a subcomplex difference operator. Given a CW complex G and subcomplex $H \subset G$, let $G \setminus H$ denote the subcomplex of G which is the closure of $G - H$, alternatively $G \setminus H$ is the union of those closed cells whose interiors are in $G - H$.

A *marked graph* is a graph G having no vertices of valence 1 equipped with a path metric and a homotopy equivalence $\rho: G \rightarrow R_n$. The fundamental group $\pi_1(G)$ is therefore identified with F_n up to inner automorphism. There is a natural identification between the

set of circuits in G , the set of conjugacy classes of $\pi_1(G)$, and the set \mathcal{C} . Each marked graph G determines a unique “length function” in $\mathbf{R}^{\mathcal{C}}$ which assigns to an element of \mathcal{C} the length of the corresponding circuit of G . Two marked graphs $\rho: G \rightarrow R_n$, $\rho': G' \rightarrow R'_n$ are equivalent if there is a homeomorphism $\theta: G \rightarrow G'$ that multiplies path metric by a constant such that $\rho'\theta$ is homotopic to ρ , and this happens if and only if the two corresponding elements of $P\mathbf{R}^{\mathcal{C}}$ are equal.

Henceforth for any marked graph G the group of deck transformations of the universal covering space \tilde{G} will be identified by F_n , under suitable choice of base points (different choice of base points only changes the action by an inner automorphism of F_n). As in Section 1.1.1 there is an induced F_n -equivariant homeomorphism between ∂F_n and the Cantor set of ends $\partial\tilde{G}$.

For any marked graph G and any subgraph $H \subset G$, the fundamental groups of the noncontractible components of H determine a free factor system in F_n that (by a moderate abuse of notation) we denote $[\pi_1 H]$ and that has one component for each noncontractible component of the subgraph H . A subgraph of G is called a *core graph* if it has no valence 1 vertices. Every subgraph $H \subset G$ has a unique core subgraph $\text{core}(H) \subset H$ which is a deformation retract of the union of the noncontractible components of H , implying that $[\pi_1 \text{core}(H)] = [\pi_1 H]$. Conversely any free factor system \mathcal{F} can be realized as $[\pi_1 H]$ for some nontrivial core subgraph H of some marked graph G .

A *path* in a graph K is a locally injective, continuous map $\gamma: J \rightarrow K$ defined on a closed, connected, nonempty subset $J \subset \mathbf{R}$ such that either γ is *trivial* meaning that J is a point, or γ can be expressed by concatenating oriented edges at vertices. Such a concatenation expression of γ , called an *edge path* representing γ , can be written in one of the following forms: a *finite path* is a finite concatenation $E_0 E_1 \dots E_n$ for $n \geq 0$; a *positive ray* is a singly infinite concatenation $E_0 E_1 E_2 \dots$; a *negative ray* is a singly infinite concatenation $\dots E_{-2} E_{-1} E_0$; and a *line* is a bi-infinite concatenation $\dots E_{-2} E_{-1} E_0 E_1 E_2 \dots$. In each case we have $E_i \neq \overline{E}_{i+1}$ by the locally injective requirement. For any path γ let $\bar{\gamma}$ or γ^{-1} denote γ with its orientation reversed; note in particular that the reversal of a positive ray is a negative ray. We regard two paths as equivalent if they have the same concatenation expression up to translation of the index in the case of a line, and up to inversion in the case of a finite path or a line, or what is the same if they differ up to a homeomorphism of their domains.

Given a continuous function $\gamma: J \rightarrow K$ where J, K are as above, and where γ takes each endpoint of J to a vertex of K , under certain conditions γ can be *tightened* to a path denoted $[\gamma]$. This is always possible when J is compact, because γ is homotopic rel endpoints either to a nontrivial path $[\gamma]$ or to a constant path in which case $[\gamma]$ denotes the corresponding trivial path. If J is noncompact then the straightened path $[\gamma]$ is defined under the additional assumption that a lift to the universal cover $\tilde{\gamma}: J \rightarrow \tilde{K}$ is a quasi-isometric embedding, in which case $[\gamma]$ is the unique path which lifts to a path in \tilde{K} that is properly homotopic to $\tilde{\gamma}$, where “properness” of the homotopy means here that the inverse image of a bounded diameter set is compact; equivalently, $[\gamma]$ is the unique path in K such that γ and $[\gamma]$ have lifts to \tilde{K} with the same finite endpoints and the same infinite ends.

A *circuit* in K is a continuous, locally injective map of an oriented circle into K . As with paths, we will not distinguish between two circuits that differ only by a homeomorphism of domains. Any circuit can be written as a concatenation of edges which is unique up

to cyclic permutation and inversion. Every continuous, homotopically nontrivial function $c: S^1 \rightarrow K$ may be tightened, meaning in this case that c is freely homotopic to a unique circuit denoted $[c]$.

A circuit represented by a local injection $\gamma: S^1 \rightarrow K$ is *root-free* if γ does not factor through a covering map $S^1 \mapsto S^1$ of positive degree. We also use the root-free terminology for a closed path $\gamma: J \rightarrow K$ to mean that γ is not a concatenation of more than one copy of the same closed path.

A path or circuit *crosses* or *contains* an edge E if E or \overline{E} occurs in its edge path expression.

1.1.5 Lines and rays.

Lines. Given any finite graph K let $\widehat{\mathcal{B}}(K)$ be the compact space of circuits and paths in K , where this space is given the *weak topology* with one basis element $\widehat{V}(K, \alpha)$ for each finite path α in K , consisting of all paths and circuits having α as a subpath. Let $\mathcal{B}(K) \subset \widehat{\mathcal{B}}(K)$ be the compact subspace consisting of all lines in K , with basis elements $V(K, \alpha) = \widehat{V}(K, \alpha) \cap \mathcal{B}(K)$.

Following [BFH00] we next turn to a coordinate free description of lines. We give the definitions in the context of an arbitrary free group A of finite rank ≥ 2 , in order to apply it to subgroups of F_n . Let $\widetilde{\mathcal{B}}(A)$ denote the set of two-element subsets of ∂A , equivalently

$$\widetilde{\mathcal{B}}(A) = (\partial A \times \partial A - \Delta)/(\mathbb{Z}/2)$$

where Δ is the diagonal and $\mathbb{Z}/2$ acts by interchanging factors. We put the *weak topology* on $\widetilde{\mathcal{B}}(A)$, induced from the usual Cantor topology on ∂A . The group A acts on $\widetilde{\mathcal{B}}(A)$ with compact but non-Hausdorff quotient space $\mathcal{B}(A) = \widetilde{\mathcal{B}}(A)/A$; we also refer to this quotient topology as the *weak topology*. A *line* of A is, by definition, an element of $\mathcal{B}(A)$. A *lift* of a line $\gamma \in \mathcal{B}(A)$ is an element $\tilde{\gamma} \in \widetilde{\mathcal{B}}(A)$ that projects to γ , and the two elements of ∂A that determine $\tilde{\gamma}$ are called its *endpoints*.

Every nontrivial conjugacy class $[b]$ in A determines a well-defined line $\gamma \in \mathcal{B}(A)$ called the *axis* of $[b]$, characterized by saying that the lifts of the axis are precisely those $\tilde{\gamma} \in \widetilde{\mathcal{B}}(A)$ which are fixed by representatives of $[b]$. Note that two nontrivial conjugacy classes $[b]$, $[b']$ have the same axis if and only if there exists $a \in A$ and integers $i, j \neq 0$ such that a^i, a^j are conjugate to b, b' respectively. The set of axes therefore forms a subset of $\mathcal{B}(A)$ in natural 1–1 correspondence with the set of unordered pairs $\{[b], [b^{-1}]\}$ where $b \in A$ is *root free* meaning that $b = c^k$ implies $k = \pm 1$.

The natural action of $\text{Aut}(A)$ on ∂A induces an action of $\text{Aut}(A)$ on $\widetilde{\mathcal{B}}(A)$ and hence an action of $\text{Out}(A)$ on $\mathcal{B}(A)$.

We also use $\widetilde{\mathcal{B}}$ and \mathcal{B} as shorthand for the spaces $\widetilde{\mathcal{B}}(F_n)$ and $\mathcal{B}(F_n)$ associated to the free group F_n .

Given any marked graph G , its space of lines $\mathcal{B}(G)$ is naturally homeomorphic to $\mathcal{B}(F_n)$, via the natural identification between the ∂F_n and the space of ends of the universal covering \widetilde{G} . Given $\gamma \in \mathcal{B}(F_n)$ the corresponding element of $\mathcal{B}(G)$ is called the *realization* of γ in G . Keeping this homeomorphism in mind, we use the term “line” ambiguously, to refer either to an element of $\mathcal{B}(F_n)$ or an element of $\mathcal{B}(G)$ in a context where a marked

graph G is under discussion. Sometimes to allay confusion an element of $\mathcal{B}(F_n)$ is called an *abstract line*.

Notation and terminology. We usually abbreviate $\tilde{\mathcal{B}}(F_n)$ to $\tilde{\mathcal{B}}$ and $\mathcal{B}(F_n)$ to \mathcal{B} , although for emphasis we sometimes do not abbreviate. We use the adjective “weak” to refer to concepts associated to the weak topology, although in contexts where this is clear we sometimes drop the adjective. In particular, although $\widehat{\mathcal{B}}(G)$ is not Hausdorff and limits of sequences are not unique, we use the phrase “weak limit” of a sequence to refer to any point to which that sequence weakly converges. As a special case, in any context where $\phi \in \text{Out}(F_n)$ or $f: G \rightarrow G$ are given, if $\alpha, \beta \in \mathcal{B}$ and $\phi^k(\alpha)$ converges weakly to β as $k \rightarrow +\infty$, or if $\alpha, \beta \in \widehat{\mathcal{B}}(G)$ and $f_\#^k(\alpha)$ converges weakly to β , then we say that α is *weakly attracted to β* (under iteration by ϕ or by $f_\#$).

Rays. A *ray of F_n* is by definition an element of the orbit set $\partial F_n/F_n$ (we will have no occasion to use the quotient topology on this set). Rays of F_n may be realized in a marked graph G as follows. Two (positive) rays in G are *asymptotic* if they have equal subrays, and this is an equivalence relation on the set of rays in G . The set of asymptotic equivalence classes of rays ρ in G is in natural bijection with $\partial F_n/F_n$ where ρ in G corresponds to end $\xi \in \partial F_n/F_n$ if there is a lift $\tilde{\rho} \subset \widetilde{G}$ of ρ , and a lift $\tilde{\xi} \in \partial F_n$ of ξ , such that $\tilde{\rho}$ converges to $\tilde{\xi}$ in the Gromov compactification of \widetilde{G} . A ray ρ in G corresponding to an abstract ray ξ is said to be a *realization of ξ in G* . Again we sometimes use the term *abstract ray* to distinguish a ray of F_n from a ray in G .

The *weak accumulation set* of a ray ρ in G is the set of lines $\ell \in \mathcal{B}(G)$ which are elements of the weak closure of ρ in the space $\widehat{\mathcal{B}}(G)$; equivalently, every finite subpath of ℓ occurs infinitely often as a subpath of ρ . Closely related is the *accumulation set* of an abstract ray $\xi \in \partial F_n/F_n$, defined as the set of lines $\ell \in \mathcal{B}$ which are in the weak closure of every line $\mu \in \mathcal{B}$ that has a lift $\tilde{\mu}$ with an endpoint $\tilde{\xi}$ projecting to ξ . We may also speak of the accumulation set of any element of ∂F_n , well-defined to be the accumulation set of its projection in $\partial F_n/F_n$. The accumulation set of a ray ρ in G is well-defined in its asymptotic equivalence class and is equal to the accumulation set of the corresponding abstract ray $\xi \in \partial F_n/F_n$.

1.1.6 The path maps $f_\#$ and $f_{\#\#}$, and bounded cancellation.

Consider a π_1 -injective map $f: K \rightarrow G$ between two finite graphs which we fix for the moment.

There is an induced continuous map $f_\#: \widehat{\mathcal{B}}(K) \rightarrow \widehat{\mathcal{B}}(G)$ defined for each path or circuit γ in K by $f_\#(\gamma) = [f(\gamma)]$, that is, $f_\#(\gamma)$ is obtained from the f -image of γ by straightening. If f is a homotopy equivalence then the restricted map $f_\#: \mathcal{B}(K) \rightarrow \mathcal{B}(G)$ is a homeomorphism. If K, G are both marked graphs then, with respect to the identifications $\mathcal{B}(K) \approx \mathcal{B} \approx \mathcal{B}(G)$, the homeomorphism $\mathcal{B} \approx \mathcal{B}(K) \xrightarrow{f_\#} \mathcal{B}(G) \approx \mathcal{B}$ equals the self-homeomorphism of \mathcal{B} induced by the outer automorphism associated to f , and in particular if f preserves marking then this self-homeomorphism is the identity.

Also define the map $f_{\#\#}: \widehat{\mathcal{B}}(K) \rightarrow \widehat{\mathcal{B}}(G)$ as follows. Intuitively $f_{\#\#}(\gamma)$ is the largest common subpath of all paths $f_\#(\delta)$ as δ varies over all paths containing γ as a subpath. To be precise, choose a lifted path $\tilde{\gamma}$ and a lifted map $\tilde{f}: \widetilde{K} \rightarrow \widetilde{G}$, define the path $\tilde{f}_{\#\#}(\tilde{\gamma}) \subset \widetilde{G}$

to be the intersection of all paths $\tilde{f}_\#(\tilde{\delta})$ as $\tilde{\delta}$ ranges over all paths in \tilde{K} that contain $\tilde{\gamma}$ as a subpath, and define $f_{\#\#}(\gamma)$ to be the projection of $\tilde{f}_{\#\#}(\tilde{\gamma})$, which is well-defined independent of the choice of $\tilde{\gamma}$. It may happen that $f_{\#\#}(\gamma)$ is empty, although by Lemma 1.6 (1) below this happens only when $f_\#(\gamma)$ is short.

Fact 1.5 (Bounded Cancellation Lemma [Coo87], [BFH97]). *For any π_1 -injective map $f: K \rightarrow G$ of finite graphs there exists a constant $\text{BCC}(f)$, called a bounded cancellation constant for f , such that for any lift $\tilde{f}: \tilde{K} \rightarrow \tilde{G}$ to universal covers and any path γ in \tilde{G} , the path $[f_\#(\gamma)]$ is contained in the $\text{BCC}(f)$ neighborhood of the image $f(\gamma)$. \square*

The following lemma will be used in the proof of Lemma 1.61 and extensively in Part II [HM13b] and Part III [HM13c]. Recall the notation $\hat{V}(\gamma, G)$ for the basis element of the weak topology on $\hat{\mathcal{B}}(G)$ associated to a finite path γ .

Lemma 1.6. *For any π_1 -injective map $f: K \rightarrow G$ of finite graphs, and for any finite path γ in K , the following hold:*

- (1) *The path $f_{\#\#}(\gamma)$ is obtained from the path $f_\#(\gamma)$ by truncating initial and terminal segments of length at most $\text{BCC}(f)$.*
- (2) $f_\#(\hat{V}(\gamma; K)) \subset \hat{V}(f_{\#\#}(\gamma); G)$.

Furthermore, for any π_1 -injective maps $f: G_1 \rightarrow G_2$ and $g: G_2 \rightarrow G_3$ of finite graphs the following hold:

- (3) *If α is a subpath of β in G_1 then $f_{\#\#}(\alpha)$ is a subpath of $f_{\#\#}(\beta)$ in G_2 .*
- (4) *If β is a path in G_1 then $g_{\#\#}(f_{\#\#}(\beta))$ is a subpath of $(g \circ f)_{\#\#}(\beta)$.*
- (5) *$f_{\#\#}$ preserves order in the following sense: if σ is a path in G_1 that decomposes into (possibly trivial) subpaths as $\sigma = \alpha_1 \beta_1 \alpha_2 \beta_2 \dots \beta_m \alpha_{m+1}$ then there is a decomposition*

$$f_\#(\sigma) = \gamma_1 f_{\#\#}(\beta_1) \gamma_2 f_{\#\#}(\beta_2) \dots f_{\#\#}(\beta_m) \gamma_{m+1}$$

for some (possibly trivial) subpaths γ_i . A similar statement holds for singly infinite and for bi-infinite decompositions of σ .

- (6) *If α is a subpath of a circuit $\sigma \subset G_1$ then $f_{\#\#}(\alpha)$ is a subpath of $f_\#(\sigma)$.*

Proof. Item (1) follows from the Bounded Cancellation Lemma. Item (2) says that if δ is a path containing γ as a subpath in K then $f_\#(\delta)$ contains $f_{\#\#}(\gamma)$ as a subpath in G ; this is just a rewriting of the definition of $f_{\#\#}$. Item (3) is an immediate consequence of the definitions. To prove (4), if β is a subpath of γ then $f_{\#\#}(\beta)$ is a subpath of $f_{\#\#}(\gamma)$ which is a subpath of $f_\#(\gamma)$, and similarly $g_{\#\#}(f_{\#\#}(\beta))$ is a subpath of $g_\#(f_\#(\gamma)) = (g \circ f)_\#(\gamma)$; since this is true for all such γ , (4) follows.

For (5), choose a lift $\tilde{\sigma} = \tilde{\alpha}_1 \tilde{\beta}_1 \tilde{\alpha}_2 \tilde{\beta}_2 \dots \tilde{\beta}_m \tilde{\alpha}_{m+1}$. Given $i < j$, write $\tilde{\sigma} = \tilde{\sigma}_1 \tilde{\sigma}_2$ where $\tilde{\sigma}_1$ is the initial subpath of $\tilde{\sigma}$ that terminates with $\tilde{\beta}_i$. It is an immediate consequence of the definition of $f_{\#\#}$ that $\tilde{f}_\#(\tilde{\sigma}_1)$ contains $\tilde{f}_{\#\#}(\tilde{\beta}_i)$ and that $\tilde{f}(\tilde{\sigma}_1)$ is disjoint from the interior of $\tilde{f}_{\#\#}(\tilde{\beta}_j)$. Similarly $f_\#(\tilde{\sigma}_2)$ contains $f_{\#\#}(\tilde{\beta}_j)$ and $\tilde{f}(\tilde{\sigma}_2)$ is disjoint from the interior of

$\tilde{f}_{\# \#}(\tilde{\beta}_i)$. It follows that $\tilde{f}_{\# \#}(\tilde{\beta}_i)$ and $\tilde{f}_{\# \#}(\tilde{\beta}_j)$ are subpaths of $\tilde{f}_{\#}(\tilde{\sigma})$ with disjoint interiors and that former precedes the latter. Since $i < j$ are arbitrary this proves (5).

To prove (6) write σ as a concatenation $\alpha\beta$ and let $\tau = f_{\#}(\sigma)$. Any line $\tilde{\sigma}$ in the universal cover \tilde{G}_1 lifting σ can be written as a bi-infinite concatenation $\tilde{\sigma} = \dots \tilde{\alpha}_{-1} \tilde{\beta}_{-1} \tilde{\alpha}_0 \tilde{\beta}_0 \tilde{\alpha}_1 \tilde{\beta}_1 \dots$ of lifts of α and β . Let T be the covering transformation that increases indices by 1, taking $\tilde{\alpha}_i$ to $\tilde{\alpha}_{i+1}$ etc. By (5) and equivariance of the $\# \#$ construction, the line $\tilde{f}_{\#}(\tilde{\sigma})$ in \tilde{G}_2 decomposes as a bi-infinite concatenation

$$\tilde{f}_{\#}(\tilde{\sigma}) = \dots \tilde{f}_{\# \#}(\tilde{\alpha}_{-1}) \tilde{\gamma}_{-1} \tilde{f}_{\# \#}(\tilde{\alpha}_0) \tilde{\gamma}_0 \tilde{f}_{\# \#}(\tilde{\alpha}_1) \tilde{\gamma}_1 \dots$$

and T increases indices by 1. It follows that $f_{\#}(\sigma) = f_{\# \#}(\alpha)\gamma$ as desired. \square

1.2 Subgroup systems carrying lines and other things.

Given a set of lines $\Lambda \subset \mathcal{B}$, in [BFH00] Section 2.6 one finds the basic definitions and facts regarding free factor systems carrying Λ , and regarding the free factor support of Λ which means the free factor system \mathcal{F} which carries Λ and which is the smallest such free factor system with respect to the partial ordering \sqsubset .

We shall need to generalize these concepts in two ways: allowing general subgroup systems; and allowing more general sets of objects beyond sets of lines. In all cases we use the operator $\mathcal{F}_{\text{supp}}(\cdot)$ to denote the free factor support.

1.2.1 Subgroup systems carrying lines and rays.

For any a finite rank subgroup $K < F_n$, the induced K -equivariant inclusion $\partial K \subset \partial F_n$ induces in turn maps of rays $\partial K/K \mapsto \partial F_n/F_n$ and of lines $\mathcal{B}(K) \mapsto \mathcal{B}$. Note that the images of these two maps depend only on the conjugacy class of the subgroup K in the group F_n .

Consider a subgroup system \mathcal{K} , a ray $\xi \in \partial F_n/F_n$, and a line $\ell \in \mathcal{B}$. We say that \mathcal{K} carries ξ if there exists a subgroup $K < F_n$ such that $[K] \in \mathcal{K}$ and ξ is in the image of the induced map $\partial K/K \mapsto \partial F_n/F_n$. We define \mathcal{K} carries ℓ similarly, requiring instead that ℓ is in the image of the induced map $\mathcal{B}(K) \mapsto \mathcal{B} = \mathcal{B}(F_n)$. Let $\mathcal{B}(\mathcal{K}) \subset \mathcal{B}$ denote the subset all lines carried by \mathcal{K} , equivalently the union of the images of the maps $\mathcal{B}(K) \mapsto \mathcal{B}$ taken over all components $[K] \in \mathcal{K}$. We do not introduce any special notation for the subset of rays carried by \mathcal{K} .

Marked graphs can be used to detect carrying of lines and rays by free factor systems. This is an evident consequence of the definitions, as noted already for lines in [BFH00] preceding Lemma 2.6.4:

Fact 1.7. *Consider a free factor system \mathcal{F} . For any marked graph G and subgraph H with $[H] = \mathcal{F}$ the following hold:*

- (1) *A line $\ell \in \mathcal{B}$ is carried by \mathcal{F} if and only if the realization of ℓ in G is contained in H .*
- (2) *An ray $\eta \in \partial F_n/F_n$ is carried by \mathcal{F} if and only if for any ray ρ in G realizing η , some subray of ρ is contained in H .* \square

Given a conjugacy class $[g]$ of F_n we say that \mathcal{K} carries $[g]$ if \mathcal{K} carries the axis of $[g]$; by Fact 1.2 this is equivalent to the existence of a subgroup $K \in \mathcal{K}$ such that $g^m \in K$ for some $m \geq 1$. In the special case that $\mathcal{K} = \mathcal{F}$ is a free factor system we have: \mathcal{F} carries a conjugacy class $[g]$ if and only if $g \in A$ for some $A \in \mathcal{F}$; and \mathcal{F} carries a line ℓ if and only if ℓ is a weak limit of a sequence of conjugacy classes carried by \mathcal{F} , if and only if for some (any) marked graph G and a subgraph $H \subset G$ with $[\pi_1 H] = \mathcal{F}$, the realization of ℓ in G is contained in H .

The following facts about subgroups carrying lines will be useful in several places.

Fact 1.8.

- (1) For any finite rank subgroup $K < F_n$ the set $\tilde{\mathcal{B}}(K)$ is closed in $\tilde{\mathcal{B}}$ and the set $\mathcal{B}(K)$ is closed in \mathcal{B} .
- (2) For any subgroup system \mathcal{K} the collection of subsets $\{\tilde{\mathcal{B}}(K) \mid K \in \mathcal{K}\}$ is uniformly locally finite in $\tilde{\mathcal{B}}$: there is a constant C and an F_n -equivariant open cover \mathcal{U} of $\tilde{\mathcal{B}}$ such that each element of \mathcal{U} intersects at most C of the subsets $\tilde{\mathcal{B}}(K)$.
- (3) For any subgroup system \mathcal{K} consisting of subgroups $\{K_i\}_{i \in I}$, and for any subset $J \subset I$, the set $\cup_{i \in J} \tilde{\mathcal{B}}(K_i)$ is closed in $\tilde{\mathcal{B}}$.

Proof. Item (1) is an immediate consequence of the fact that ∂K is a compact subset of ∂F_n and the fact that the action of K on its minimal subtree in \tilde{R}_n is cocompact. Items (1) and (2) together clearly imply (3).

To prove (2) we work in the Cayley graph \tilde{R}_n . For any finite rank subgroup $A < F_n$ let $T(A) \subset \tilde{R}_n$ be the minimal subtree for the action of A on \tilde{R}_n , which equals the convex hull in \tilde{R}_n of the Gromov boundary $\partial A \subset \partial R_n = \partial F_n$. For each edge $e \subset \tilde{R}_n$ let $\tilde{\mathcal{B}}(e) \subset \mathcal{B}$ be the open subset consisting of lines whose realization in \tilde{R}_n contains e , and let $\mathcal{U} = \{\tilde{\mathcal{B}}(e)\}$. Pick subgroups $K_1, \dots, K_L < F_n$ such that $\mathcal{K} = \{[K_1], \dots, [K_L]\}$. For each $l = 1, \dots, L$ and each edge e we shall bound the size of any subset $\Gamma \subset F_n$ representing pairwise distinct left cosets of K_l such that if $\gamma \in \Gamma$ then $e \subset \gamma(T(K_l)) = T(\gamma K_l \gamma^{-1})$, equivalently $\tilde{\mathcal{B}}(e) \cap \tilde{\mathcal{B}}(\gamma K_l \gamma^{-1}) \neq \emptyset$. Pick a finite subtree $D_l \subset T(K_l)$ which is a fundamental domain for the action of K_l on $T(K_l)$. For each $\gamma \in \Gamma$, after replacing γ by γk for appropriate $k \in K_l$ we may assume that $e \subset \gamma D_l$, and by proper discontinuity of the action of F_n on \tilde{R}_n the set Γ has bounded cardinality $C_{l,e}$ depending only on l and e . Noticing that $C_{l,e}$ depends only on the orbit of e under the action of F_n , by restricting to edges e in a fundamental domain for the action of F_n on \tilde{R}_n we obtain a finite maximum $C = \max C_{l,e}$, finishing the proof. \square

1.2.2 Free factor supports of lines, rays, and subgroup systems.

The *free factor support* of a subset of lines $B \subset \mathcal{B}$, denoted $\mathcal{F}_{\text{supp}}(B)$, is defined by the following result from [BFH00]:

Fact 1.9 ([BFH00], Corollary 2.6.5). *For any subset $B \subset \mathcal{B}$ the meet of all free factor systems that carry B is a free factor system that carries B that is denoted $\mathcal{F}_{\text{supp}}(B)$. If B is a single line then $\mathcal{F}_{\text{supp}}(B)$ has one component. $\mathcal{F}_{\text{supp}}(B)$ is the unique free factor system carrying B that is minimal with respect to the partial ordering \sqsubset .*

Fact 1.9 is also sufficient to incorporate the *free factor support* of a set of conjugacy classes C , denoted $\mathcal{F}_{\text{supp}}(C)$, by which we simply mean the free factor support of the set of axes for C .

We need to define the free factor support $\mathcal{F}_{\text{supp}}(X)$ of a polyglot set X each of whose elements is a line, conjugacy class, ray, or subgroup system of F_n . This is done using the following generalization of Fact 1.9:

Fact 1.10. *For any set X each of whose elements is either a line, conjugacy class, ray, or subgroup system of F_n , the following hold:*

- (1) *The meet of all free factor systems that carry each element of X is the unique free factor system that carries each element of X and is minimal with respect to \sqsubset . This free factor system is called the free factor support of X , denoted $\mathcal{F}_{\text{supp}}(X)$.*
- (2) *$\mathcal{F}_{\text{supp}}(X)$ is natural in the sense that for each $\phi \in \text{Out}(F_n)$ we have $\mathcal{F}_{\text{supp}}(\phi(X)) = \phi(\mathcal{F}_{\text{supp}}(X))$. In particular if X is ϕ -invariant then $\mathcal{F}_{\text{supp}}(X)$ is ϕ -invariant.*
- (3) *If X is a single line or ray then $\mathcal{F}_{\text{supp}}(X)$ has one component.*

Proof. For the proof of (1), as in the proof of Fact 1.9 given in [BFH00] Corollary 2.6.5, because the meet of an arbitrary collection of free factor systems equals the meet of some finite subcollection, and because meet is associative, it suffices to consider just two free factor systems $\mathcal{F}_1, \mathcal{F}_2$, and to prove that if a line, conjugacy class, ray, or subgroup system is carried by \mathcal{F}_1 and by \mathcal{F}_2 then it is carried by $\mathcal{F}_1 \wedge \mathcal{F}_2$. For conjugacy classes the definition of carrying is a special case of the definition for lines and so need not be considered further. Recall the proof for a line ℓ carried by $\mathcal{F}_1, \mathcal{F}_2$: choosing a lift $\tilde{\ell} \in \tilde{\mathcal{B}}$, there exist $A_1, A_2 < F_n$ such that $[A_i] \in \mathcal{F}_i$ and $\partial\ell \subset \partial A_i$, and so $\partial\ell \subset \partial A_1 \cap \partial A_2 = \partial(A_1 \cap A_2)$ (see Fact 1.2); by Fact 1.3 we have $[A_1 \cap A_2] \in \mathcal{F}_1 \wedge \mathcal{F}_2$, so ℓ is carried by $\mathcal{F}_1 \wedge \mathcal{F}_2$. Similarly, for a ray $\xi \in X$ carried by $\mathcal{F}_1, \mathcal{F}_2$, choosing a lift $\tilde{\xi} \in \partial F_n$, again there exist A_i as above with $\tilde{\xi} \in \partial A_i$ and so $\tilde{\xi} \in \partial(A_1 \cap A_2)$ and so ξ is carried by $\mathcal{F}_1 \wedge \mathcal{F}_2$. For a subgroup system $\mathcal{K} \sqsubset F_n$, it suffices to consider each component $[K] \in \mathcal{K}$ one at a time. For $i = 1, 2$ there exists $A_1, A_2 < F_n$ such that $K < A_i$, and so $K < A_1 \cap A_2$ implying that $[K] \sqsubset \mathcal{F}_1 \wedge \mathcal{F}_2$.

The naturally clause (2) is evident. The proof of (3) when X is a single line is given shown in [BFH00], Corollary 2.6.5, and the same proof works if X is a single ray ρ , using the evident fact that if a free factor system \mathcal{F} carries ρ then some component of \mathcal{F} carries ρ . \square

The following fact relates the free factor support of a subgroup system \mathcal{K} to the free factor support of the set of lines carried by \mathcal{K} :

Fact 1.11. *Given a subgroup system \mathcal{K} the following hold:*

- (1) *For any free factor system \mathcal{F} , the inclusion $\mathcal{K} \sqsubset \mathcal{F}$ holds if and only if each line carried by \mathcal{K} is carried by \mathcal{F} .*
- (2) *$\mathcal{F}_{\text{supp}}(\mathcal{B}(\mathcal{K})) = \mathcal{F}_{\text{supp}}(\mathcal{K})$.*

Proof. The implication (1) \implies (2) is obvious, as is the “only if” direction of (1).

For the “if” direction of (1), given $K < F_n$ so that $[K] \in \mathcal{K}$, it suffices to assume that $\mathcal{B}(K)$ is carried by \mathcal{F} and that $[K] \not\sqsubset \mathcal{F}$, and to derive a contradiction. Choose a nontrivial $\gamma_1 \in K$. Choose A_1 so that $[A_1] \in \mathcal{F}$ and $\partial\gamma_1 \subset \partial A_1$. Since $[K] \not\sqsubset \mathcal{F}$ it

follows that $K \not\prec A_1$. Choosing $\gamma_2 \in K - A_1$ it follows that $\partial\gamma_2 \not\subset \partial A_1$, for otherwise $\partial A_1 \cap \partial(\gamma_2 A_1 \gamma_2^{-1}) = \partial A_1 \cap \gamma_2(\partial A_1) \neq \emptyset$ contradicting malnormality of A_1 and Fact 1.2. Choose $A_2 \in \mathcal{F}$ such that $\partial\gamma_2 \subset \partial A_2$. Clearly $\partial\gamma_1 \neq \partial\gamma_2$ and $A_1 \neq A_2$. Pick $\xi_1 \in \partial\gamma_1$ and $\xi_2 \in \partial\gamma_2$ such that $\xi_1 \neq \xi_2$, let $\tilde{\ell} = \{\xi_1, \xi_2\} \in \mathcal{B}$, and let $\ell \in \mathcal{B}$ be its projection. By construction we have $\ell \in \tilde{\mathcal{B}}(K)$ so we may choose $A_3 \in \mathcal{F}$ such that $\partial\tilde{\ell} \subset \partial A_3$. It follows that $\partial A_1 \cap \partial A_3 \neq \emptyset$ and $\partial A_2 \cap \partial A_3 \neq \emptyset$. But since one of A_1 or A_2 is not equal to A_3 this contradicts malnormality of \mathcal{F} and Fact 1.2. \square

Remark. By similar proofs we can show that the equivalent statements in Fact 1.11 (1) are also equivalent to the statement that each ray carried by \mathcal{K} is carried by \mathcal{F} , and that the subgroup system in (2) is also equal to the free factor support of the set of rays carried by \mathcal{K} . We do not need these results so we omit the proofs.

The following fact gives a bit more useful information relating free factor systems to the lines and rays that they carry.

Fact 1.12. *For any free factor system \mathcal{F} ,*

- (1) *For every sequence of lines $\ell_i \in \mathcal{B}$, if every weak limit of every subsequence of ℓ_i is carried by \mathcal{F} then ℓ_i is carried by \mathcal{F} for all sufficiently large i .*
- (2) *The weak accumulation set of every ray not carried by \mathcal{F} contains a line not carried by \mathcal{F} .*
- (3) *The free factor support of every ray is equal to the free factor support of its weak accumulation set.*

Proof. If the conclusion of (1) fails then some subsequence ℓ_{i_n} is not carried by \mathcal{F} . Choose a marked graph G and a core subgraph H representing \mathcal{F} . The realization of ℓ_{i_n} in G is not contained in H and so, after passing to a subsequence, each ℓ_{i_n} contains some edge $E \subset G - H$. It follows that some weak limit of some subsequence is a line that contains E , and so is not contained in H and is not carried by \mathcal{F} .

To prove (2), if the ray r in G realizes an abstract ray not carried by \mathcal{F} then there is an edge $E \subset G - H$ that the ray r crosses infinitely many times, and we can proceed as in (1).

To prove (3), given a ray $\xi \in \partial F_n/F_n$ with weak accumulation set $\Lambda \subset \mathcal{B}$, the inclusion $\mathcal{F}_{\text{supp}}(\xi) \subset \mathcal{F}_{\text{supp}}(\Lambda)$ is a consequence of (2), and the reverse inclusion $\mathcal{F}_{\text{supp}}(\Lambda) \subset \mathcal{F}_{\text{supp}}(\xi)$ follows by Fact 1.7 (2). \square

1.3 Attracting laminations.

We recall here the basic definitions and properties of attracting laminations proved in Section 3 of [BFH00].

A closed subset of $\mathcal{B} = \mathcal{B}(F_n)$ is called a *lamination*, or an F_n -*lamination* when emphasis is needed. An element of a lamination is called a *leaf*. The action of $\text{Out}(F_n)$ on \mathcal{B} induces an action on the set of laminations.

For each marked graph G the homeomorphism $\mathcal{B} \approx \mathcal{B}(G)$ induces a bijection between F_n -laminations and closed subsets of $\mathcal{B}(G)$. The closed subset of $\mathcal{B}(G)$ corresponding to a lamination $\Lambda \subset \mathcal{B}$ is called the *realization* of Λ in G ; we often conflate Λ with its realizations

in marked graphs. Also, we occasionally use the term lamination to refer to F_n -invariant, closed subsets of $\tilde{\mathcal{B}}$; the orbit map $\tilde{\mathcal{B}} \mapsto \mathcal{B}$ puts these in natural bijection with laminations in \mathcal{B} .

A line $\ell \in \mathcal{B}$ is *birecurrent* if ℓ is in the weak accumulation set of every subray of ℓ . A lamination Λ is *minimal* if each of its leaves is dense in Λ . Every leaf of a minimal lamination is birecurrent, and a connected lamination in which every leaf is birecurrent is minimal.

Definition 1.13 (Attracting laminations). Given $\phi \in \text{Out}(F_n)$ and a lamination $\Lambda \subset \mathcal{B}$, we say that Λ is an *attracting lamination* for ϕ if there exists a leaf $\ell \in \Lambda$ satisfying the following: Λ is the weak closure of ℓ ; ℓ is a birecurrent; ℓ is not the axis of the conjugacy class of a generator of a rank one free factor of F_n ; and there exists $p \geq 1$ and a weak open set $U \subset \mathcal{B}$ such that $\phi^p(U) \subset U$ and such that $\{\phi^{kp}(U) \mid k \geq 1\}$ is a weak neighborhood basis of ℓ . Any such leaf ℓ is called a *generic leaf* of Λ . Let $\mathcal{L}(\phi)$ be the set of attracting laminations for ϕ .

Fact 1.14. *For each $\phi \in \text{Out}(F_n)$ the following hold:*

- (1) $\mathcal{L}(\phi)$ is finite and ϕ -invariant.
- (2) For all $\Lambda \in \mathcal{L}(\phi)$ and each generic leaf ℓ of Λ we have $\mathcal{F}_{\text{supp}}(\Lambda) = \mathcal{F}_{\text{supp}}(\ell)$, and this free factor system has one component.
- (3) For all $\Lambda, \Lambda' \in \mathcal{L}(\phi)$ we have $\Lambda = \Lambda'$ if and only if $\mathcal{F}_{\text{supp}}(\Lambda) = \mathcal{F}_{\text{supp}}(\Lambda')$.
- (4) There is a bijection $\mathcal{L}(\phi) \leftrightarrow \mathcal{L}(\phi^{-1})$ defined by $\Lambda_+ \leftrightarrow \Lambda_-$ if and only if $\mathcal{F}_{\text{supp}}(\Lambda_+) = \mathcal{F}_{\text{supp}}(\Lambda_-)$. \square

Definition 1.15 (Dual pairs of laminations). The bijection defined in item (4) is called *duality* between $\mathcal{L}(\phi)$ and $\mathcal{L}(\phi^{-1})$, a pair $\Lambda^\pm = (\Lambda^+, \Lambda^-)$ corresponding in this manner is called a *dual pair* of laminations of ϕ , and the set of dual pairs is denoted $\mathcal{L}^\pm(\phi)$.

Remark. Given $\Lambda \in \mathcal{L}(\phi)$, if $\phi(\Lambda) = \Lambda$ then in Definition 1.13 we may assume $p = 1$, because if $p > 1$ then by Fact 1.14 (1) we may replace U by $U \cap \phi(U) \cap \dots \cap \phi^{p-1}(U)$. In general we may assume that p is the period of Λ , meaning the least integer for which $\phi(p) = \Lambda$.

Remark. In general the action of ϕ on the finite set $\mathcal{L}(\phi)$ may be a nontrivial permutation. But in certain circumstances one sees that each element of $\mathcal{L}(\phi)$ is fixed: this holds for $\phi \in \text{IA}_n(\mathbb{Z}/3)$ by Lemma II.4.2¹, and for rotationless ϕ (see Definition 1.22 below) by Lemma 3.30 of [FH11].

1.4 Principal automorphisms and rotationless outer automorphisms.

The normal subgroup $\text{Inn}(F_n)$ acts by conjugation on $\text{Aut}(F_n)$, and the orbits of this action define an equivalence relation on $\text{Aut}(F_n)$ called *isogredience*, which is a refinement of the relation of being in the same outer automorphism class. In this section we shall define various isogredience invariants of elements of $\text{Aut}(F_n)$.

¹ “Lemma II.X.Y” or “Section II.V.W” refers to Lemma X.Y or Section V.W of Part II [HM13b].

Notation 1.16. One inevitably abbreviates ‘forward rotationless’ to ‘rotationless’ so we will do that from the start. Throughout this series of papers, forward rotationless outer automorphisms are referred to as rotationless outer automorphisms.

Definition 1.17. The group $\text{Aut}(F_n)$ acts on ∂F_n . Let $\widehat{\Phi}: \partial F_n \rightarrow \partial F_n$ denote the continuous extension of the action of $\Phi \in \text{Aut}(F_n)$, and $\text{Fix}(\widehat{\Phi}) \subset \partial F_n$ its set of fixed points. Let $\text{Fix}(\Phi) < F_n$ denote the subgroup of elements fixed by Φ , shown first to be of finite rank in [Coo87] and then to be of rank $\leq n$ in [BH92]. Its boundary $\partial \text{Fix}(\Phi) \subset \partial F_n$ is empty when $\text{Fix}(\Phi)$ is trivial, 2 points when $\text{Fix}(\Phi)$ has rank 1, and a Cantor set when $\text{Fix}(\Phi)$ has rank ≥ 2 . Let $\text{Fix}_+(\widehat{\Phi})$ denote the set of attractors in $\text{Fix}(\widehat{\Phi})$, a discrete subset consisting of points $\xi \in \text{Fix}(\widehat{\Phi})$ such that for some neighborhood $U \subset \partial F_n$ of ξ we have $\widehat{\Phi}(U) \subset U$ and the sequence $\widehat{\Phi}^n(\eta)$ converges to ξ for each $\eta \in U$. Let $\text{Fix}_-(\widehat{\Phi}) = \text{Fix}_+(\widehat{\Phi}^{-1})$ denote the set of repellers in $\text{Fix}(\widehat{\Phi})$. For a nontrivial $\gamma \in F_n$ the associated homeomorphism at infinity, denoted $\widehat{\gamma}: \partial F_n \rightarrow \partial F_n$, is identical to continuous extension of the action of the associated inner automorphism $i_\gamma \in \text{Inn}(F_n)$, and we denote the 2-point fixed set as $\text{Fix}(\widehat{\gamma}) = \{\text{Fix}_-(\widehat{\gamma}), \text{Fix}_+(\widehat{\gamma})\}$. The following lemma combines results of [GJLL98] and [BFH04]; see [FH11] Lemma 2.1 for detailed citations.

Fact 1.18. *For any $\Phi \in \text{Aut}(F_n)$, and for any nontrivial $\gamma \in F_n$ with the following are equivalent:*

- (1) $\text{Fix}(\widehat{\gamma}) \cap \text{Fix}(\widehat{\Phi}) \neq \emptyset$
- (2) $\text{Fix}(\widehat{\gamma}) \subset \text{Fix}(\widehat{\Phi})$
- (3) $\gamma \in \text{Fix}(\Phi)$
- (4) Φ commutes with i_γ .

□

The following simple corollary of Fact 1.18 will be used often without comment:

Fact 1.19. *If $\Phi_1 \neq \Phi_2 \in \text{Aut}(F_n)$ represent the same outer automorphism then $\text{Fix}(\widehat{\Phi}_1) \cap \text{Fix}(\widehat{\Phi}_2)$ is either empty or equal to $\text{Fix}(\widehat{\gamma})$ where $\Phi_1^{-1}\Phi_2 = i_\gamma$.* □

In the following statement, items (1) and (2) follow immediately from [GJLL98], Proposition I.1.

Fact 1.20. *For each $\Phi \in \text{Aut}(F_n)$ we have*

$$\text{Fix}(\widehat{\Phi}) = \partial \text{Fix}(\Phi) \cup \text{Fix}_-(\widehat{\Phi}) \cup \text{Fix}_+(\widehat{\Phi}) \quad (*)$$

Also, if $\text{Fix}(\widehat{\Phi}) \neq \text{Fix}(\widehat{\gamma})$ for each nontrivial $\gamma \in F_n$ then the following hold:

- (1) *The expression (*) is a disjoint union.*
- (2) *Each $\xi \in \text{Fix}_+(\widehat{\Phi})$ is an attracting point for the action of Φ on the Gromov compactification $F_n \cup \partial F_n$.*
- (3) *Each term on both sides of (*) is invariant under the action of $\text{Fix}(\Phi)$ on ∂F_n . The points of $\text{Fix}_+(\widehat{\Phi})$ are the isolated points of the set $\partial \text{Fix}(\Phi) \cup \text{Fix}_+(\widehat{\Phi})$; the same holds with $\text{Fix}_-(\widehat{\Phi}) = \text{Fix}_+(\widehat{\Phi}^{-1})$ in place of $\text{Fix}_+(\widehat{\Phi})$.*

Proof. We prove (3). Given $\gamma \in \text{Fix}(\Phi)$, by Fact 1.18 the actions of $\widehat{\Phi}$ and $\widehat{\gamma}$ on ∂F_n commute; invariance of the terms in $(*)$ under $\widehat{\gamma}$ follows immediately. In the set $\partial \text{Fix}(\Phi) \cup \text{Fix}_+(\widehat{\Phi})$, isolation of points of $\text{Fix}_+(\widehat{\Phi})$ is evident, and nonisolation of points of $\text{Fix}(\widehat{\Phi})$ is evident when $\text{rank}(\text{Fix}(\Phi)) \neq 1$ because $\text{Fix}(\widehat{\Phi})$ is either empty or a Cantor set. When $\text{rank}(\text{Fix}(\Phi)) = 1$ it follows by assumption that $\text{Fix}_+(\widehat{\Phi}) \neq \emptyset$, and clearly each $\text{Fix}(\Phi)$ orbit of a point in $\text{Fix}_+(\widehat{\Phi})$ accumulates on both points of $\partial \text{Fix}(\Phi)$. \square

Let $\text{Fix}_N(\widehat{\Phi}) = \text{Fix}(\widehat{\Phi}) - \text{Fix}_-(\widehat{\Phi})$, which equals $\partial \text{Fix}(\Phi) \cup \text{Fix}_+(\widehat{\Phi})$ as long as $\text{Fix}(\widehat{\Phi})$ is not the fixed point pair of a nontrivial element of F_n . Let $\text{Per}(\widehat{\Phi}) = \cup_{k \geq 1} \text{Fix}(\widehat{\Phi}^k)$, and similarly for $\text{Per}_+(\widehat{\Phi})$, $\text{Per}_-(\widehat{\Phi})$, and $\text{Per}_N(\widehat{\Phi})$.

Fact 1.21 ([FH11] Lemmas 3.23 and 3.26, or [LL08] Theorem I). *For every $\Phi \in \text{Aut}(F_n)$ there exists $m > 0$ such that $\text{Fix}_N(\widehat{\Phi}^m) \neq \emptyset$.*

Remark. Our notation of $\text{Fix}_+(\Phi) \subset \partial F$ is not quite the same as the notion in [GJLL98] of an “attracting fixed point” in ∂F_n under the combined action of $\Phi \cup \widehat{\Phi}$ on the Gromov compactification $F_n \cup \partial F_n$, but the difference can occur only in exceptional situation described in Fact 1.20. There are indeed examples in this situation, such as automorphisms of the form $i_\alpha^k \Phi'$ where $\Phi' \in \text{Aut}(F_n)$, $\alpha \in \text{Fix}(\Phi')$, and k is sufficiently large. Nonetheless, no point of $\partial \text{Fix}(\Phi)$ is ever an “attracting fixed point” in the sense of [GJLL98], because each point of $\partial \text{Fix}(\Phi)$ is a limit in $F_n \cup \partial F_n$ of elements of F_n fixed by Φ . We note that the exceptional situation of Fact 1.20 (??) occurs for principal automorphisms (see Definition 1.22).

The following combines Definitions 3.1 and 3.13 of [FH11]:

Definition 1.22 (Principal automorphisms and rotationless outer automorphisms.). We say that $\Phi \in \text{Aut}(F_n)$ is a *principal automorphism* if $\text{Fix}_N(\widehat{\Phi})$ contains at least three points, or $\text{Fix}_N(\widehat{\Phi})$ contains exactly two points which are neither the endpoints of some axis nor the endpoints of a lift of a generic leaf of an element of $\mathcal{L}(\phi)$. Note that if Φ is principal then by Fact 1.20 we have a disjoint union $\text{Fix}_N(\widehat{\Phi}) = \partial \text{Fix}(\Phi) \cup \text{Fix}_+(\widehat{\Phi})$. The set of principal automorphisms representing $\phi \in \text{Out}(F_n)$ is denoted $P(\phi)$, it is invariant under isogredience, and by [FH11] Remark 3.9 there are only finitely many isogredience classes.

We say that $\phi \in \text{Out}(F_n)$ is *(forward) rotationless* if $\text{Fix}_N(\widehat{\Phi}) = \text{Per}_N(\widehat{\Phi})$ for all $\Phi \in P(\phi)$, and if for each $k \geq 1$ the map $\Phi \mapsto \Phi^k$ induces a bijection between $P(\phi)$ and $P(\phi^k)$.

Fact 1.23 ([FH11] Lemma 4.42). *There exists K depending only on the rank n such that for each $\phi \in \text{Out}(F_n)$, ϕ^K is rotationless.* \square

Fact 1.24 ([FH11] Lemma 3.30 and Corollary 3.31). *If $\phi \in \text{Out}(F_n)$ is rotationless then*

- (1) *Every conjugacy class in F_n which is ϕ -periodic is fixed by ϕ .*
- (2) *ϕ fixes every element of $\mathcal{L}(\phi)$.*
- (3) *Every free factor system in F_n which is ϕ -periodic is fixed by ϕ .*
- (4) *If F is a ϕ -invariant free factor then $\phi \mid F$ is rotationless.* \square

1.5 Relative train track maps and CTs

In this section we review general relative train track maps, as well as the more specialized CTs or completely split relative train track maps.

1.5.1 Definitions of relative train track maps and CTs

Topological representatives and Nielsen paths. Given $\phi \in \text{Out}(F_n)$ a *topological representative* of ϕ is a map $f: G \rightarrow G$ such that $\rho: R_n \rightarrow G$ is a marked graph, f is a homotopy equivalence, f takes vertices to vertices and edges to paths, and $\bar{\rho} \circ f \circ \rho: R_n \rightarrow R_n$ represents ϕ .

A nontrivial path γ in G is a *periodic Nielsen path* if there exists k such that $f_\#^k(\gamma) = \gamma$. The minimal such k is the *period*, and if $k = 1$ then γ is a *Nielsen path*. A periodic Nielsen path is *indivisible* if it cannot be written as a concatenation of two nontrivial periodic Nielsen paths. We will often abuse our notion of path equivalence when talking about Nielsen paths: many statements asserting the uniqueness of an indivisible Nielsen path ρ , such as Fact 1.40, should really assert uniqueness of ρ up to reversal of direction; since ρ determines its reversal $\bar{\rho}$ and vice versa, this abuse is easily detectable and harmless.

Filtrations. A *filtration* of a marked graph G is a strictly increasing sequence of subgraphs $G_0 \subset G_1 \subset \dots \subset G_k = G$, each with no isolated vertices. The individual terms G_k are called *filtration elements*, and if G_k is a core graph then it is called a *core filtration element*. The subgraph $H_k = G_k \setminus G_{k-1}$ is called the *stratum of height k*. The *height* of a subset of G is the minimum k such that the subset is contained in G_k . The height of a map to G is the height of the image of the map. The height of an abstract ray $\xi \in \partial F_n / F_n$ equals s if and only if every ray in G realizing ξ has a subray contained in G_s and not contained in G_{s-1} ; equivalently, for any ray R in G realizing ξ , the highest stratum having edges crossed infinitely often by R is H_s . A *connecting path* of a stratum H_k is a nontrivial finite path γ of height $< k$ whose endpoints are contained in H_k .

Given a topological representative $f: G \rightarrow G$ of $\phi \in \text{Out}(F_n)$, we say that f respects the filtration or that the filtration is *f-invariant* if $f(G_k) \subset G_k$ for all k . If this is the case then we also say that the filtration is *reduced* if for each free factor system \mathcal{A} which is invariant under ϕ^i for some $i \geq 1$, if $[\pi_1 G_{r-1}] \sqsubset \mathcal{A} \sqsubset [\pi_1 G_r]$ then either $\mathcal{A} = [\pi_1 G_{r-1}]$ or $\mathcal{A} = [\pi_1 G_r]$.

Given an *f*-invariant filtration, for each stratum H_k with edges $\{E_1, \dots, E_m\}$, define the *transition matrix* of H_k to be the square matrix whose i, j entry equals the number of times that the edge E_j occurs in the path $f(E_i)$. If M_k is the zero matrix then we say that H_k is a *zero stratum*. If M_k irreducible — meaning that for each i, j there exists p such that the i, j entry of the p^{th} power of the matrix is nonzero — then we say that H_k is *irreducible*; and if one can furthermore choose p independently of i, j then H_k is *aperiodic*. Assuming that H_k is irreducible, by [Haw08] the matrix M_k a unique eigenvalue $\lambda \geq 1$, called the *Perron-Frobenius eigenvalue*, for which some associated eigenvector has positive entries: if $\lambda > 1$ then we say that H_k is an *exponentially growing* or EG stratum; whereas if $\lambda = 1$ then H_k is a *nonexponentially growing* or NEG stratum.

Directions and turns. A *direction* of a marked graph G at a point $x \in G$ is a germ of finite paths with initial vertex x . If x is not a vertex then the number of directions at

x equals 2. If x is a vertex then the number of directions equals the valence of x , in fact there is a natural bijection between the directions at x and the oriented edges with initial vertex x , and we shall often identify a direction at x with its corresponding oriented edge. Let $T_x G$ be the set of directions of G at x and let TG be the union of $T_x G$ over all vertices x of G . A *turn* at $x \in G$ is an unordered pair of directions $\{d, d'\} \subset T_x G$. If $d \neq d'$ then the turn is *nondegenerate*, otherwise it is *degenerate*.

If a filtration of G is given, the *height* of a direction d is well-defined as the height of any sufficiently short path representing d , equivalently the height of the oriented edge representing d . A nondegenerate turn is said to have height r if each of its directions has height r .

Given a marked graph G and a homotopy equivalence $f: G \rightarrow G$ that takes edges to paths, a turn $\{d, d'\}$ in G is said to be *legal* with respect to the action of f if the turn $\{(Df)^k(d), (Df)^k(d')\}$ is nondegenerate for all $k \geq 0$. For any path γ and any turn $\{E, E'\}$ at a vertex, if $\bar{E}E'$ or its inverse $\bar{E}'E$ is a subpath of γ then we say that the turn $\{E, E'\}$ is *taken* by γ . If in addition an f -invariant filtration is given, a path γ is r -*legal* if it has height $\leq r$ and each turn of height r taken by γ is legal.

Definition 1.25. Relative train track maps. We define relative train track maps using the method of [FH11] Section 2.6, which imposes certain mild preconditions on a filtered topological representative, and which explains how an arbitrarily filtered topological representative can be mildly altered, by subdivision of edges and/or refinement of the filtration, so as to satisfy these preconditions.

Consider $\phi \in \text{Out}(F_n)$ and a topological representative $f: G \rightarrow G$ of ϕ with f -invariant filtration $G_0 \subset G_1 \subset \dots \subset G_k = G$ satisfying the following preconditions:

- (1) Every stratum is either a zero stratum or irreducible.
- (2) The edges of each NEG stratum H_k may be assigned orientations called *NEG orientations*, and they may be numbered as E_1, \dots, E_N , so that for all $i \in \mathbb{Z}/N\mathbb{Z}$ we have $f(E_i) = E_{i+1} u_i$ where $u_i \subset G_{k-1}$ is a (possibly empty) path.

Note that while item (2) is stable under passage to a positive power $f^k: G \rightarrow G$, item (1) need not be. However there always exists *some* positive power so that, after further refinement of the invariant filtration, each irreducible stratum is aperiodic, which forces item (1) to become stable under passage to any further positive power.

We say that f is a *relative train track* representative of ϕ if for each EG stratum H_r the following hold:

RTT-(i) Df maps the set of directions of height r to itself. In particular, each turn consisting of a direction of height r and one of height $< r$ is legal.

RTT-(ii) For each connecting path γ of H_r , $f_\#(\gamma)$ is a connecting path of H_r .

RTT-(iii) For each r -legal path α , the path $f_\#(\alpha)$ is r -legal.

For each NEG stratum H_k we use the following additional terminology. If each u_i is the trivial path then we say that H_k is a *periodic stratum* and each edge $E \subset H_k$ is a *periodic edge*; if furthermore $N = 1$ then H_k is a *fixed stratum* and E is a *fixed edge*. If each u_i is

either trivial or a periodic Nielsen path, with at least one periodic Nielsen path, then we say that H_k is a *linear stratum* and each $E \subset H_k$ is a *linear edge*. An NEG edge E which is neither a fixed edge nor a linear edge is said to be a *superlinear edge*. This completes the definition of a relative train track map.

For the most part we will be interested in the special class of relative train track maps called CTs, recounted in Definition 1.29 after some further preliminaries.

Definition 1.26. Principal vertices. Consider a relative train track map $f: G \rightarrow G$ with filtration $\emptyset = G_0 \subset G_1 \subset \dots \subset G_N$. Two periodic points $x \neq y \in G$ are *Nielsen equivalent* if there exists a periodic Nielsen path with endpoints x, y . A periodic point x is *principal* if and only if *neither* of the following hold:

- (1) x is the unique periodic point in its Nielsen class, there are exactly two periodic directions based at x , and both of those directions are in the same EG stratum.
- (2) x is contained in a topological circle $C \subset G$ such that each point of C is periodic and has exactly two periodic directions (but see Remark following Fact 1.37 which says that no such periodic point x exists in the case that f is a CT).

Definition 1.27. Splittings. Let $f: G \rightarrow G$ be a relative train track map with invariant filtration $\emptyset = G_0 \subset G_1 \subset \dots \subset G_N = G$. A *splitting* of a path or circuit σ in G is a decomposition of σ into a concatenation of subpaths $\sigma = \sigma_1 \cdot \dots \cdot \sigma_k$ such that for all $i \geq 1$ the path or circuit $f_\#^i(\sigma)$ decomposes as $f_\#^i(\sigma_1) \dots f_\#^i(\sigma_k)$. The single dot notation “.” indicates that the decomposition is a splitting, and the subpaths $\sigma_1, \dots, \sigma_k$ are called the *terms* of the splitting. In other words, a decomposition of σ into a concatenation of subpaths is a splitting of σ if one can tighten the image of σ under any iterate by simply tightening the images of the subpaths.

Certain paths will be atomic terms of splittings, including the following two types.

Given a zero stratum H_i , a path τ in H_i , and an irreducible stratum H_j with $j > i$, we say that τ is *j-taken* if there exists an edge $E \in H_j$ and $k \geq 1$ such that τ is a maximal subpath in H_i of the path $f_\#^k(E)$. We also say that τ is *taken* if it is *j-taken* for some j .

Given two linear edges E_i, E_j , a root-free closed Nielsen path w , and integers d_i, d_j of the same sign such that $f_\#(E_i) = E_i w^{d_i}$ and $f_\#(E_j) = E_j w^{d_j}$, each path of the form $E_i w^p \bar{E}_j$ for $p \in \mathbb{Z}$ is called an *exceptional path*.

Definition 1.28. Complete Splittings. A splitting of a path or circuit $\sigma = \sigma_1 \cdot \dots \cdot \sigma_k$ is called a *complete splitting* if each term σ_i satisfies one of the following: σ_i is an edge in an irreducible stratum; σ_i is an indivisible Nielsen path; σ_i is an exceptional path; there is a zero stratum H_j such that σ_i is a maximal subpath of σ in H_j , and σ_i is taken. We say that σ is *completely split* if it has a complete splitting.

Enveloping of zero strata. Consider a relative train track map $f: G \rightarrow G$ with filtration $\emptyset = G_0 \subset G_1 \subset \dots \subset G_N = G$ and two strata H_u, H_r with $1 \leq u < r \leq N$. Suppose that the following hold:

- (1) H_u is irreducible.

- (2) H_r is EG and each component of G_r is noncontractible.
- (3) Each H_i with $u < i < r$ is a zero stratum that is a component of G_{r-1} , and each vertex of H_i has valence ≥ 2 in G_r .

In this case we say that the zero strata H_i with $u < i < r$ are *enveloped* by the EG stratum H_r , and we write $H_r^z = \cup_{i=u+1}^r H_i$.

Definition 1.29 (CT — Completely split relative train track map. [FH11] Definition 4.7). A CT is a relative train track map with particularly nice properties. Of the nine defining properties, all but the ninth involve terms that are fully defined in this article. The statement of the ninth, which we shall not ever use, involves terms for definitions of which we refer the reader to [FH11].

Consider a relative train track map $f: G \rightarrow G$ with filtration $\emptyset = G_0 \subset G_1 \subset \dots \subset G_K = G$. The following properties define what it means for f to be a CT:

- (1) (**Rotationless**) Each principal vertex is fixed by f and each periodic direction at a principal vertex is fixed by Df (c.f. [FH11] Definition 3.18).
- (2) (**Completely Split**) For each edge E in each irreducible stratum, the path $f(E)$ is completely split. For each taken connecting path σ in each zero stratum, the path $f_\#(\sigma)$ is completely split.
- (3) (**Filtration**) The filtration $\emptyset = G_0 \subset G_1 \subset \dots \subset G_N = G$ is reduced. For each i there exists $j \leq i$ such that $G_i = \text{core}(G_j)$.
- (4) (**Vertices**) The endpoints of all indivisible Nielsen paths are vertices. The terminal endpoint of each nonfixed NEG edge is principal.
- (5) (**Periodic Edges**) Each periodic edge is fixed and each endpoint of a fixed edge is principal. If $E = H_r$ is a fixed edge and E is not a loop then G_{r-1} is a core graph and both ends of E are contained in G_{r-1} .
- (6) (**Zero Strata**) Each zero stratum H_i is enveloped by some EG stratum H_r , each edge in H_i is r -taken, and each vertex in H_i is contained in H_r and has link contained in $H_i \cup H_r$.
- (7) (**Linear Edges**) For each linear edge E_i there exists a closed root-free Nielsen path w_i such that $f(E_i) = E_i w_i^{d_i}$ for some $d_i \neq 0$. If E_i, E_j are distinct linear edges, and if the closed path w_i is freely homotopic to the closed path w_j or its inverse, then $w_i = w_j$ and $d_i \neq d_j$.
- (8) (**NEG Nielsen Paths**) If $H_i = \{E_i\}$ is an NEG stratum and σ is an indivisible Nielsen path of height i then E_i is linear and, with w_i as in (Linear Edges), there exists $k \neq 0$ such that $\sigma = E_i w_i^k \overline{E}_i$.
- (9) (**EG Nielsen Paths**) If H_i is an EG stratum and ρ is an indivisible Nielsen path of height i then $f|_{G_i} = \theta \circ f_{i-1} \circ f_i$ where: the map $f_i: G_i \rightarrow G^1$ is a composition of proper extended folds defined by iteratively folding ρ ; the map $f_{i-1}: G^1 \rightarrow G^2$ is a composition of folds involving edges in G_{i-1} ; and $\theta: G^2 \rightarrow G_i$ is a homeomorphism.

The main existence theorem for CTs is the following. Given a nested sequence $\mathcal{C} = (\mathcal{F}_1 \sqsubset \cdots \sqsubset \mathcal{F}_L)$ of free factor systems and a topological representative $f: G \rightarrow G$ with f -invariant filtration $\emptyset = G_0 \subset G_1 \subset \cdots \subset G_K = G$, we say that f realizes \mathcal{C} if for each $\ell = 1, \dots, L$ there exists $k_\ell \in 1, \dots, K$ such that G_{k_ℓ} is a core subgraph that realizes \mathcal{F}_ℓ .

Theorem 1.30 ([FH11] Theorem 4.28). *For each rotationless $\phi \in \text{Out}(F_n)$ and each nested sequence \mathcal{C} of ϕ -invariant free factor systems, there exists a CT $f: G \rightarrow G$ that represents ϕ and realizes \mathcal{C} . \square*

1.5.2 Facts about CTs, their Nielsen paths, and their zero strata.

For Facts 1.31–1.45, let $f: G \rightarrow G$ be a CT with filtration $\emptyset = G_0 \subset G_1 \subset \cdots \subset G_k$, representing a rotationless $\phi \in \text{Out}(F_n)$.

Fact 1.31. *If H_r is an EG stratum then H_r is aperiodic and G_r is a core graph.*

Proof. By [FH11] Lemma 3.19, H_r contains a principal vertex v whose link contains a principal direction in H_r . By (Rotationless), that direction is fixed. The transition matrix of H_r therefore has a nonzero element on the diagonal, which implies that H_r is aperiodic. By (Filtration) and [FH11] Lemma 2.20 (2), G_r is a core graph. \square

Fact 1.32 ([FH11], Remark 4.9). *A vertex of G whose link contains edges in more than one irreducible stratum is principal.* \square

Fact 1.33 ([FH11], Lemma 4.11). *Every completely split path or circuit has a unique complete splitting.* \square

The next fact is used repeatedly without reference:

Fact 1.34 ([FH11], Lemma 4.6). *If σ is a path in G with endpoints at vertices and if σ is completely split then $f_\#(\sigma)$ is completely split. Moreover, if $\sigma = \sigma_1 \dots \sigma_k$ is a complete splitting then $f_\#(\sigma)$ has a complete splitting that refines $f_\#(\sigma) = f_\#(\sigma_1) \dots f_\#(\sigma_k)$.* \square

Fact 1.35. *If σ is a circuit in G or a path in G with endpoints at vertices then $f_\#^k(\sigma)$ is completely split for all sufficiently large $k \geq 0$.*

Proof. When σ is a path this is [FH11] Lemma 4.25. When σ is a circuit then, by [BFH00] Lemma 4.1.2, σ may be written as a based closed path σ_1 which splits at its base point, and so the path $f_\#^k(\sigma_1)$ completely splits for sufficiently large k . By [BFH00] Lemma 4.1.1 any complete splitting of the path $f_\#^k(\sigma_1)$ gives a splitting of the circuit $f_\#^k(\sigma)$, which is evidently a complete splitting of the circuit. \square

Fact 1.36 ([FH11], Lemma 4.21). *Each nonfixed NEG stratum H_i is a single edge E_i which, with its NEG orientation, has a splitting $f(E_i) = E_i \cdot u_i$ where u_i is a nontrivial, closed, completely split circuit of height $< i$.* \square

Fact 1.37. *For each periodic vertex $x \in G$, exactly one of the following holds:*

- (1) *x is principal (and therefore fixed);*
- (2) *x is the unique periodic point in its Nielsen class, there are exactly two periodic directions based at x , and both of those directions are in the same EG stratum.*

Proof. If x is contained in a topological circle C of periodic points then, applying (Periodic Edges), each edge of C is fixed and (1) holds. If x is contained in no such circle then exactly one of items (1) or (2) holds, by definition of principal. \square

Remark: The above argument shows that circles as Definition 1.26 (2), the definition of principal vertices, do not exist in CT's.

Nielsen paths.

Fact 1.38 ([FH11] Lemma 4.13). *Every periodic Nielsen path is a (fixed) Nielsen path.*

Fact 1.39. *If σ is an $f_{\#}$ -periodic path or circuit then σ is $f_{\#}$ -fixed, σ is completely split, and all terms in the complete splitting of σ are fixed edges or indivisible Nielsen paths.*

Proof. Suppose σ is $f_{\#}$ periodic, say $f_{\#}^N(\sigma) = \sigma$. By Fact 1.35, $\sigma = f_{\#}^{kN}(\sigma)$ is completely split for sufficiently large k . By Fact 1.34, applied to $\sigma = f_{\#}^N(\sigma)$, all terms in the complete splitting of σ are $f_{\#}$ -periodic Nielsen paths. By Fact 1.38, all terms are Nielsen paths, and so σ is $f_{\#}$ -fixed. \square

Our next few facts are concerned with indivisible Nielsen paths of EG height.

Fact 1.40. *For each EG stratum H_r , up to reversal there is at most one indivisible Nielsen path ρ_r of height r ([FH11] Corollary 4.19). Moreover, if ρ_r exists then ρ_r may be written uniquely as a concatenation $\rho = \alpha\beta$ where α and β are r -legal paths, each begins and ends with edges of H_r , and the turn $\{Df(\bar{\alpha}), Df(\beta)\}$ is degenerate ([BH92] Lemma 5.11).*

Fact 1.41. *If H_r is an EG stratum, if there exists an indivisible Nielsen path ρ_r of height r , and if $\rho_r = a_0b_1a_1b_2 \cdots a_{k-1}b_k$ is the decomposition into maximal subpaths a_i in H_r and b_i in G_{r-1} , then:*

- (1) *No zero stratum is enveloped by H_r .*
- (2) *Each b_i is a Nielsen path.*
- (3) *For each edge $E \subset H_r$ and each $k \geq 0$, the path $f_{\#}^k(E)$ splits into terms each of which is an edge of H_r or one of the Nielsen paths b_i in the decomposition of ρ_r . Furthermore, each term in the complete splitting of $f_{\#}^k(E)$ is an edge of H_r , a fixed edge, or an indivisible Nielsen path.*

Proof. Lemma 4.24 of [FH11] implies (1), (2) and the first statement of (3) for the case $k = 1$; the case $k > 1$ follows by induction. The furthermore part of (3) follows from Corollary 4.12 of [FH11]. \square

Fact 1.42. *If H_r is an EG stratum, and if ρ_r is an indivisible Nielsen path of height r , then ρ_r crosses each edge of H_r at least once, the initial oriented edges of ρ_r and $\bar{\rho}_r$ are distinct oriented edges of H_r , and:*

(1) *ρ_r is not closed if and only if it crosses some edge of H_r exactly once, and in this case:*

- (a) *At least one endpoint of ρ_r is not in G_{r-1} .*
- (b) *There does not exist a height r fixed conjugacy class. In particular, there does not exist a closed, height r Nielsen path.*
- (c) *There exists a proper free factor system \mathcal{F} such that for each line $\ell \in \mathcal{B}$, ℓ is carried by \mathcal{F} if and only if the realization of ℓ in G is a concatenation of edges of $G \setminus H_r$ and copies of ρ_r .*

(2) *ρ_r is closed if and only if it crosses each edge of H_r exactly twice, and in this case:*

- (a) *The endpoint of ρ_r is not in G_{r-1} .*
- (b) *The only height r fixed conjugacy classes are those represented by ρ_r , its inverse, and their iterates.*

Proof. The statement about initial oriented edges is found in [FH11] Corollary 4.19 eg-(i), which derives from [BFH00] Theorem 5.1.5 eg-(i). The statements regarding numbers of edge crossings of H_r made by ρ_r are found in [BH92] Theorem 5.15 but with a hypothesis saying “ f is stable”, but all that is needed for the proof of Theorem 5.15 to go through are the properties of folds contained in (EG Nielsen Paths); see also [BFH00] Lemma 5.2.5 for the same statements but with a hypothesis saying that “ f is \mathcal{F} -Nielsen minimized”.

Item (1c) is just Lemma 5.1.7 of [BFH00].

If ρ_r crosses some edge of H_r exactly once then the hypotheses of [BFH00] Lemma 5.1.7 hold, the conclusion of which says that ρ_r is not closed and that at most one of its endpoints is in a noncontractible component of G_{r-1} , but by Facts 1.41 and 1.44, every component of G_{r-1} is noncontractible. If ρ_r crosses every edge of H_r exactly twice, it follows by [BFH00] Proposition 5.3.1 that ρ_r is a closed path which begins and ends with distinct edges and that its endpoint is not contained in G_{r-1} .

It remains to prove items (1b) and (2b). Consider a height r circuit c fixed by $f_\#$. By Fact 1.35, c completely splits. Each term of the complete splitting must be a fixed edge of height $< r$ or an indivisible Nielsen path of height $\leq r$, including at least one term of height r which, by Fact 1.40, must be ρ_r or its inverse. Up to switching orientations of c and of ρ_r , we may assume that $c_i = \rho_r$ is a term of σ with terminal endpoint $x \notin G_{r-1}$. Consider the term c_{i+1} following c_i , with initial endpoint x . Since no fixed edge in G_r is incident to x , c_{i+1} cannot be a fixed edge. If ρ_r is not closed then the only candidate for c_{i+1} is the inverse of ρ_r , which is not allowed to follow ρ_r in a complete splitting. If ρ_r is closed then the only allowed candidate for c_{i+1} is another copy of ρ_r , and by induction c is an iterate of ρ_r . \square

Our last fact about Nielsen paths is concerned with zero strata and with NEG strata that are *superlinear*, meaning not fixed and not linear.

Fact 1.43. *If H_r is a zero stratum or an NEG superlinear stratum then no Nielsen path crosses any edge of H_r .*

Proof. Suppose some Nielsen path crosses an edge E of H_r , and let i be the minimal height of all such paths. By Fact 1.39, there is an indivisible Nielsen path ρ_i of height i that crosses E_r . Note that H_i is note a zero stratum for otherwise, by (Zero Strata), it would follows that ρ_i is contained entirely in H_i , but $f^k(H_i) \cap H_i = \emptyset$ for sufficiently large k , contradicting that $f_\#^k(\rho_i) = \rho_i$.

Case 1: $i = r$. We have shown that $H_i = H_r$ is not a zero stratum, and if H_r is NEG superlinear then by applying (NEG Nielsen Paths) we obtain a contradiction.

Case 2: $i > r$. If $H_i = \{E_i\}$ is NEG then by (NEG Nielsen Paths) it follows that H_i is linear and that $\rho_i = E_i w_i^s \overline{E}_i$ where w_i is a closed Nielsen path of height $< i$, but then w_i crosses E_r , contradicting minimality of height. If H_i is EG then by Fact 1.41 we have $\rho_i = a_0 b_1 a_1 b_2 \cdots a_{k-1} b_k$ where the a 's are paths in H_i and the b 's are Nielsen paths of height $< i$, but then one of the b 's must cross E_r , again contradicting minimality of height. We have already ruled out the possibility that H_i is a zero stratum. \square

Zero strata.

Fact 1.44. *For each filtration element G_r the following are equivalent:*

- (1) G_r has a contractible component;
- (2) H_r is a zero stratum;
- (3) H_r is a contractible component of G_r .

Proof. The equivalence of (1) and (2) is Lemma 4.15 of [FH11]. That (2) implies (3) is a consequence of (Zero Strata) which implies that H_r is a component of G_r , and the fact that $f_\#$ is a bijection on circuits. That (3) implies (1) is obvious. \square

Fact 1.45. *For any edge $E \subset G$, if $f_\#(E)$ is contained in a zero stratum H_t enveloped by the EG stratum H_s then E is an edge in some other zero stratum $H_{t'} \neq H_t$ that is also enveloped by H_s .*

Proof. Since the path $f_\#(E)$ is completely split and contained in the zero stratum H_t , its complete splitting must have just the single term $f_\#(E)$ which must be a taken connecting path of H_s . If the oriented edge E is contained in an irreducible stratum H_i then some term in the complete splitting of $f_\#(E)$ is an edge of H_i , contradiction. It follows that E is contained in some zero stratum $H_{t'}$ enveloped by some EG stratum $H_{s'}$.

By (Zero Strata) applied to $H_{t'}$ there exists an oriented edge $E_0 \subset H_{s'}$ having the same initial vertex as E . The paths $f_\#(E_0)$ and $f_\#(E)$ therefore have the same initial vertex, and by (Zero Strata) applied to H_t the link of this vertex is contained in $H_s \cup H_t$. By RTT-(i) the initial direction of $f_\#(E_0)$ is contained in $H_{s'}$, and so $s = s'$.

If $H_t = H_{t'}$ then $f(H_t) \cap H_t \neq \emptyset$, contradicting the definition of a zero stratum. \square

1.5.3 Facts about principal lifts, principal directions, and principal rays.

From here to the end of the section we fix a rotationless $\phi \in \text{Out}(F_n)$ represented by a CT $f: G \rightarrow G$. We describe several natural structures on f that are associated to principal automorphisms representing ϕ .

Recall the natural bijection between the set of automorphisms Φ representing ϕ , and the set of lifts $\tilde{f}: \tilde{G} \rightarrow \tilde{G}$ of $f: G \rightarrow G$ via the universal covering map $\tilde{G} \mapsto G$, where \tilde{f} corresponds to Φ if and only if the extension \tilde{f} to $\partial\tilde{G} \approx \partial F_n$ equals the extension $\widehat{\Phi}$ to ∂F_n .

Definition 1.46 ([FH11] Definition 3.1). Given a lift $\tilde{f}: \tilde{G} \rightarrow \tilde{G}$ with corresponding automorphism Φ representing ϕ , we say that \tilde{f} is a *principal lift* if Φ is a principal automorphism.

Fact 1.47 ([FH11] Corollary 3.17, Corollary 3.22, Corollary 3.27). *A lift $\tilde{f}: \tilde{G} \rightarrow \tilde{G}$ is principal if and only if there is a principal vertex $v \in G$ and a lift $\tilde{v} \in \tilde{G}$ such that $\tilde{f}(\tilde{v}) = \tilde{v}$.* \square

Definition 1.48 (Principal directions). Consider an oriented edge $E \subset G$ whose initial vertex is a principal vertex of f and whose initial direction is fixed by f but E is not a fixed edge of f . If E is nonlinear then we say that E is a *principal direction of f in G* . Given a principal lift $\tilde{f}: \tilde{G} \rightarrow \tilde{G}$, and given an oriented edge $\tilde{E} \subset \tilde{G}$ with initial vertex $\tilde{v} \in \text{Fix}(\tilde{f})$, we say that \tilde{E} is a *principal direction of \tilde{f} in \tilde{G}* if its projection $E \subset G$ is a principal direction of f in G . We also say that E (or \tilde{E}) is an NEG or EG principal direction depending on the nature of the stratum of G containing E .

In the following fact, the hypotheses on Φ and \tilde{f} are satisfied when Φ is principal, equivalently \tilde{f} is a principal lift. While most of its numerous applications are in the principal case, there are a few applications in the nonprincipal case. Following its statement we use this fact to define principal rays, then we give the proof by mostly just citing Lemmas 3.26 and 4.36 of [FH11].

Fact 1.49. *Let $\Phi \in \text{Aut}(F_n)$ represent ϕ with associated lift $\tilde{f}: \tilde{G} \rightarrow \tilde{G}$. Suppose that $\text{Fix}(\tilde{f}) \neq \emptyset$, and suppose that there does not exist a nontrivial $\gamma \in F_n$ such that the set $\text{Fix}(\widehat{\Phi}) = \text{Fix}(\tilde{f})$ is equal to $\text{Fix}(\widehat{\gamma})$. For each $\tilde{v} \in \text{Fix}(\tilde{f})$ and each oriented edge $\tilde{E} \subset \tilde{G}$ with initial vertex \tilde{v} and with fixed initial direction, let $v, E \subset G$ be the projections of \tilde{v}, \tilde{E} . With this notation, the following hold:*

- (1) *For each \tilde{v}, \tilde{E} as above, if $E = H_r$ is a fixed edge then there exists a ray $\tilde{R} \subset \tilde{G}$ with initial edge \tilde{E} that converges to some $\tilde{\xi} \in \text{Fix}_N(\widehat{\Phi})$ and that projects to a ray in G_r .*
- (2) *For each \tilde{v}, \tilde{E} as above, if E is not a fixed edge then, letting $f(E) = Eu$, letting $\tilde{u} \subset \tilde{G}$ be the lift of u for which $\tilde{f}(\tilde{E}) = \tilde{E}\tilde{u}$, and letting H_r be the stratum containing E , the following hold. There exists a properly nested sequence of oriented paths $\tilde{E} \subset \tilde{f}(\tilde{E}) \subset \tilde{f}_\#^2(\tilde{E}) \subset \dots$, each an initial segment of the previous, whose union is a ray $\tilde{R} \subset \tilde{G}$ that contains no fixed point other than \tilde{v} and that converges to some $\tilde{\xi} \in \text{Fix}_N(\widehat{\Phi})$. Furthermore:*

- (a) *There exists a splitting $f(\tilde{E}) = \tilde{E} \cdot \tilde{u}$ and an induced splitting*

$$\tilde{R} = \underbrace{\tilde{E} \cdot \tilde{u} \cdot \tilde{f}_\#(\tilde{u}) \cdot \tilde{f}_\#^2(\tilde{u}) \cdot \dots \cdot \tilde{f}_\#^k(\tilde{u}) \cdot \tilde{f}_\#^{k+1}(\tilde{u}) \cdot \dots}_{\tilde{f}_\#^k(\tilde{E})} \quad (*)$$

- (b) If E is not an NEG-linear edge then $\tilde{\xi} \in \text{Fix}_+(\widehat{\Phi})$. Furthermore, if H_r is EG then the weak accumulation set of ξ is the unique attracting lamination of height r .
 - (c) If E is an NEG-linear edge then $\tilde{\xi} \in \partial \text{Fix}(\Phi)$.
 - (d) \tilde{R} projects to a ray in G_r , and if H_r is an NEG-stratum then $\tilde{R} \setminus \tilde{E}$ projects to a ray in G_s where s is the height of the circuit u .
- (3) For each $\tilde{\xi} \in \text{Fix}_+(\widehat{\Phi})$ there exists $\tilde{v} \in \text{Fix}(\tilde{f})$ and \tilde{E} as above such that $\tilde{\xi}$ is obtained from \tilde{E} as in (2).

Definition 1.50 (Principal rays and linear rays). In the context of Fact 1.49 (2) where $E \subset G$ is a non-fixed edge with fixed initial direction, we say that \tilde{R} is *the ray in \tilde{G} generated by E* . Its projection to G is *the ray in G generated by E* which may be written in the form $R = E \cdot u \cdot f_\#(u) \cdot f_\#^2(u) \cdots$. If E is a linear edge then we say that R is a *linear ray* in G , and if E is a principal direction then R is a *principal ray*; we use similar language for \tilde{R} in \tilde{G} .

Note that if $\Phi \in P(\phi)$ with associated principal lift \tilde{f} then by definition we have a bijection between the set of principal rays of \tilde{f} in \tilde{G} and the set of principal directions of \tilde{f} in \tilde{G} , which associates to each principal ray its initial direction. Adding in Fact 1.49 items (2c) and (3) we also have a surjection from the set of principal rays of \tilde{f} in \tilde{G} to the set $\text{Fix}_+(\widehat{\Phi})$, which associates to each principal ray \tilde{R} its endpoint. Similarly, we have a bijection between principal rays of f in G and principal directions of f in G .

Remark. In Section II.2.4 we shall focus on those principal rays which are generated by NEG principal directions, and we shall show that such rays are invariants of ϕ in that they can be characterized independently of relative train tracks. Such rays, in the appropriate context, are referred to as “eigenrays” of ϕ .

Proof of Fact 1.49. Item (1) is proved in [FH11] Lemma 3.26. The main clause of item (2) is exactly [FH11] Lemma 4.36 (1), whose proof implicitly proves item (2a) as well. Also, item (3) is the same as [FH11] Lemma 4.36 (2). We note that while Lemma 4.36 assumes that \tilde{f} is principal, its proof holds under our present weaker hypotheses on \tilde{f} . Item (2d) is clear from the construction of \tilde{R} ; see also [FH11] Lemma 3.26 (2) and (3).

The “Furthermore” clause of (2b) follows from [FH11] Lemma 3.26 (2). For the rest of the proof of (2b) and (2c), E is linear if and only if $f_\#(u) = u$, and we must prove that this holds if and only if $\tilde{\xi} \notin \text{Fix}_+(\widehat{\Phi})$.

If $E \subset G$ is linear then $f_\#^k(u) = u$ for all $k \geq 1$. By construction \tilde{f} takes the initial vertex of \tilde{u} to its terminal vertex and so, letting $T_{\tilde{u}}$ be the covering transformation which does the same to those two vertices, it follows that \tilde{f} commutes with $T_{\tilde{u}}$, that $T_{\tilde{u}}(\tilde{\xi}) = \tilde{\xi}$, and that $\tilde{\xi} \in \partial \text{Fix}(\widehat{\Phi})$. It follows by Fact 1.20 that $\tilde{\xi} \notin \text{Fix}_+(\widehat{\Phi})$.

Assuming that $\tilde{\xi} \notin \text{Fix}_+(\widehat{\Phi})$, and using the hypotheses on \tilde{f} , it follows that $\tilde{\xi} \in \partial \text{Fix}(\Phi)$ by applying Fact 1.20 and Definition 1.22 (which derive from [GJLL98]). In particular the subgroup $\text{Fix}(\Phi)$ is nontrivial. The nonempty set $\text{Fix}(\tilde{f})$ is $\text{Fix}(\Phi)$ invariant and it accumulates on $\partial \text{Fix}(\Phi)$. Since the ray \tilde{R} ends at $\tilde{\xi} \in \partial \text{Fix}(\Phi)$ it follows that for $x \in \tilde{R}$ there is a uniform upper bound to $d(x, \text{Fix}(\tilde{f}))$ and so there is a uniform upper bound to $d(x, \tilde{f}(x))$. Applying (2a) it follows that there is a uniform upper bound to the length of $f_\#^k(\tilde{u})$, which implies that E is linear. \square

The following lemma will be used in Part III [HM13c].

Lemma 1.51. *For each $\Phi \in P(\phi)$ and $P \in \text{Fix}_+(\Phi)$ there is a conjugacy class $[a]$ that is weakly attracted to every line in the weak accumulation set of P .*

Proof. Let $\tilde{f}: \tilde{G} \rightarrow \tilde{G}$ be the principal lift corresponding to Φ . Applying Fact 1.49, let \tilde{R} be a principal ray with endpoint P , let \tilde{E} be its initial oriented edge, a principal direction of \tilde{f} . Consider the splitting of \tilde{R} expressed in item (2a) of Fact 1.49. Since u is not fixed by $f_\#$ we have $\text{Length}(f_\#^k(u)) \rightarrow +\infty$ as $k \rightarrow +\infty$. For sufficiently large k , say $k \geq k_0$, the paths $f_\#^k(u)$ are circuits whose initial directions are independent of k , as are their terminal directions. Let $[a]$ be the conjugacy class represented by the circuit $\sigma = f_\#^{k_0}(u) \cdot f_\#^{k_0+1}(u)$, and we have a circuit $f_\#^i(\sigma) = f_\#^{k_0+i}(u) \cdot f_\#^{k_0+i+1}(u)$.

Given a line ℓ in G which is in the accumulation set of P , and given a finite subpath α of ℓ , let U_α be the set of lines whose realization in G contains α as a subpath. Choosing a lift $\tilde{\alpha} \subset G'$, the ray \tilde{R} contains infinitely many translates of $\tilde{\alpha}$. The length of these translates are fixed while the lengths of the paths $\tilde{f}_\#^j(u)$ go to $+\infty$ with j , so infinitely many of these paths intersect translates of $\tilde{\alpha}$ in \tilde{R} , and hence infinitely many of the paths $\tilde{f}_\#^j(u) \cdot \tilde{f}_\#^{j+1}(u)$ contain translates of $\tilde{\alpha}$. It follows that infinitely many iterates $f_\#^i(\sigma)$ contain α as a subpath and so are in U_α . Since the U_α form a weak neighborhood basis of ℓ , we have proved that $f_\#^i(\sigma)$ weakly accumulates on ℓ , i.e. $[a]$ is weakly attracted to ℓ . \square

1.5.4 Properties of EG strata.

Recall Fact 1.35, which derives from [FH11] Lemma 4.25, saying that for any CT $f: G \rightarrow G$ and any σ which is either a path with endpoints at vertices or a circuit, there exists some $k \geq 1$ so that $f^k(\sigma)$ completely splits. We prove two lemmas with more quantitative control on the exponent k , achieved at the expense of discarding “completeness” of the splitting, using instead a somewhat coarser kind of splitting. The lemmas are stated for general relative train tracks—CTs are not needed—and for circuits whose height is equal to the height of some EG stratum. Their proofs derive from [BFH00] Section 4.2, although the first lemma can be derived from the EG case of the proof of [FH11] Lemma 4.25 by adding a little quantitative information.

Lemma 1.52. *Let $f: G \rightarrow G$ be a relative train track and $H_r \subset G$ an aperiodic EG stratum. For each L there exists an integer $k = k_L \geq 0$ such that for each height r path or circuit σ whose endpoints, if any, are vertices, if σ has length $\leq L$ then $f_\#^k(\sigma)$ splits into terms each of which is an edge or indivisible periodic Nielsen path of height r or a path in G_{r-1} .*

Proof. Since there are only finitely many circuits and paths with endpoints at vertices that have length $\leq L$, it suffices to consider just a single σ and find a K such that $f_\#^K(\sigma)$ satisfies the conclusions.

Let P_r be the set of paths ρ of height r such that for each $k \geq 0$ the path $f_\#^k(\rho)$ begins and ends with an edge in H_r and has exactly one r -illegal turn, and such that the number of edges in $f_\#^k(\rho)$ is bounded independently of k . By [BFH00] Lemma 4.2.5, P_r is a finite $f_\#$ invariant set, from which it follows that there exists K_1 such that for each $\rho \in P_r$ the path $f_\#^{K_1}(\rho)$ is a periodic Nielsen path, which therefore splits into indivisible periodic Nielsen paths of height r and paths in G_{r-1} .

Since $f_\#$ applied to an r -legal path is r -legal, and since a subpath of an r -legal path is r -legal, it follows that the number of maximal r -legal paths in $f_\#^i(\sigma)$ is a nonincreasing function of i . The number of illegal turns of $f_\#^i(\sigma)$ in H_r is therefore nonincreasing, and so is constant for sufficiently large i , say $i \geq K_2$. Applying Lemma 4.2.6 of [BFH00] it follows that $f_\#^{K_2}(\sigma)$ splits into subpaths each of which is r -legal or is one of the paths in P_r . Letting $K = K_1 + K_2$ it follows that $f_\#^K(\sigma)$ splits into r -legal paths and indivisible periodic Nielsen paths of height r . Since an r -legal path splits into edges in H_r and paths in G_{r-1} , this finishes the proof. \square

We also need a form of Lemma 1.52 whose conclusion has a stronger uniformity. This will be used in the analysis of limit trees in Section II.6.

Lemma 1.53. *Let $f: G \rightarrow G$ be a relative train track and $H_r \subset G$ an aperiodic EG stratum. For each M there exists an integer $d \geq 0$ such that for each height r path or circuit σ with endpoints, if any, at vertices, if σ contains at most M edges in the subgraph H_r then $f_\#^d(\sigma)$ splits into terms each of which is an edge or indivisible periodic Nielsen path of height r or a path in G_{r-1} .*

Proof. In this proof, all paths and circuits are in G and have endpoints, if any, at vertices.

If a path or circuit α of height r has a splitting satisfying the conclusions of the lemma, the terms being edges and indivisible Nielsen paths of height r and paths in G_{r-1} , then $f_\#(\alpha)$ also has a splitting satisfying the conclusions. We are therefore free to increase the exponent on $f_\#$ as needed in this proof, which we shall do without mention.

Case 1: σ is a height r path. We prove the lemma in this case by induction on M . If σ is r -legal, in particular if $M = 1$, we can take $d = 0$. Assume by induction that the exponent $D = d_{M-1}$ works when σ has $\leq M - 1$ edges in H_r .

Let $B \geq 0$ be a bounded cancellation constant for f^D . Let L be a constant such that if β is a path of length $> L$ then the path $f_\#(\beta)$ has length $\geq 2B + 1$; the constant L exists because the lift of f to the universal cover of G is a quasi-isometry. Given a path $\alpha\beta\gamma$, if β has length $> L$ then not all of $f_\#^D(\beta)$ is cancelled when $f_\#^D(\alpha)f_\#^D(\beta)f_\#^D(\gamma)$ is tightened to $f_\#^D(\alpha\beta\gamma)$: at most B initial edges of $f_\#^D(\beta)$ cancel with $f_\#^D(\alpha)$, and at most B terminal edges cancel with $f_\#^D(\gamma)$.

Let d_M be the maximum of D and the constant $K = k_{M+ML-L}$ from Lemma 1.52. Let σ be a height r path with exactly M edges in H_r .

We first reduce to the subcase that σ begins and ends with edges in H_r . In the case of a general path σ , we can write $\sigma = \alpha\tau\beta$ where α, β are the longest initial and terminal segments in G_{r-1} . Knowing that the path $f_\#^{d_M}(\tau)$ satisfies the conclusions, it has a splitting $f_\#^{d_M}(\tau) = \alpha' \cdot \tau' \cdot \beta'$ where α', β' are its longest initial and terminal segments in G_{r-1} and where τ' satisfies the conclusions, and so $f_\#^{d_M}(\sigma) = [f_\#^{d_M}(\alpha)\alpha'] \cdot \tau' \cdot [\beta' f_\#^{d_M}(\beta)]$ has a splitting that satisfies the conclusions, completing the reduction.

Suppose now that σ begins and ends with edges in H_r . If σ has length $\leq M + ML - L$ then $h_\#^K(\sigma)$ satisfies the conclusions. If σ has length $> M + ML - L$ then, since σ begins and ends with edges in H_r , we can write $\sigma = \alpha\beta\gamma$ where β in G_{r-1} has length $> L$ and α, γ

each have between 1 and $M - 1$ edges in H_r . The induction hypothesis implies that $f_\#^D(\alpha)$ and $f_\#^D(\gamma)$ each satisfy the conclusions, so there are splittings

$$\begin{aligned} f_\#^D(\alpha) &= \alpha' \cdot \beta_1 \\ f_\#^D(\gamma) &= \beta_2 \cdot \gamma' \end{aligned}$$

where β_1, β_2 are each maximal subpaths in G_{r-1} , the paths α', γ' satisfy the conclusions, and the terminal edge of α' and the initial edge of γ' are both in H_r . Our choice of L guarantees that $\beta' = [\beta_1 f_\#^D(\beta) \beta_2]$ is nontrivial, and so

$$\begin{aligned} f_\#^D(\sigma) &= [f_\#^D(\alpha) f_\#^D(\beta) f_\#^D(\gamma)] \\ &= \alpha' \beta' \gamma' \end{aligned}$$

This is a splitting of $f_\#^D(\sigma)$, because the turns $\{\bar{\alpha}', \beta'\}, \{\bar{\beta}', \gamma'\}$ are legal by RTT-(i), and α', γ' have splittings that satisfy the conclusions. It follows that $f_\#^D(\sigma)$ has a splitting that satisfies the conclusions.

Case 2: σ is a height r circuit. Suppose that σ contains exactly M edges in H_r . Using the constant d_M from Case 1, let B' be a bounded cancellation constant for f^{d_M} . Let L' be a constant such that if β is a path of length $> L'$ then the path $f_\#^{d_M}(\beta)$ has length $\geq 2B' + 1$. Let d'_M be the maximum of d_M and the constant $K' = k_{M(L'+1)}$ from Lemma 1.52. If σ has length $\leq M(L' + 1)$ then $f_\#^{K'}(\sigma)$ satisfies the conclusions. If σ has length $> M(L' + 1)$ then σ has a maximal subpath β in G_{r-1} of length $> L'$, with a complementary subpath τ that begins and ends in H_r and has exactly M edges in H_r . Applying Case 1, the path $f_\#^{d_M}(\tau)$ satisfies the conclusions, so there is a splitting

$$f_\#^{d_M}(\tau) = \beta_1 \cdot \tau' \cdot \beta_2$$

where β_1, β_2 are maximal subpaths in G_{r-1} , τ' satisfies the conclusions, and τ' begins and ends with edges of H_r . The choice of L' guarantees that $\beta' = [\beta_2 f_\#^{d_M}(\beta) \beta_1]$ is nontrivial, and so $f_\#^{d_M}(\sigma) = \beta' \tau'$. This is a splitting of the circuit $f_\#^{d_M}(\sigma)$, because $\{\bar{\beta}', \tau'\}$ and $\{\bar{\tau}', \beta'\}$ are legal by RTT-(i) and τ' has a splitting that satisfies the conclusions. It follows that $f_\#^{d_M}(\sigma)$ has a splitting that satisfies the conclusions.

In conclusion, we have proved the lemma with $d = \max\{d_M, d'_M\}$. □

1.6 Properties of Attracting Laminations

In Section 1.3 we recalled the definition of attracting laminations, which is formulated invariantly without regard to relative train tracks. We also reviewed there several facts which are formulated invariantly (albeit their proofs often depend on relative train tracks). In this section we prove various facts and lemmas about attracting laminations which are explicitly formulated from the relative train track point of view.

1.6.1 The relation between EG strata and attracting laminations.

Basic to the theory of attracting laminations is the following fact which describes the bijection between attracting laminations and EG strata. This fact is compiled from results in [BFH00] Section 3, from Definition 3.1.5 to Definition 3.1.14:

Fact 1.54. *For any $\phi \in \text{Out}(F_n)$ and any relative train track representative $f: G \rightarrow G$, ϕ acts as the identity on $\mathcal{L}(\phi)$ if and only if each EG stratum of G is EG-aperiodic. If this is the case then there exists a natural bijection between attracting laminations $\Lambda \in \mathcal{L}(\phi)$ and EG-strata $H_r \subset G$ such that Λ corresponds to H_r if and only if the following equivalent statements hold:*

- (1) *Each generic leaf of Λ has height r .*
- (2) *$\mathcal{F}_{\text{supp}}(\Lambda) \sqsubset [\pi_1 G_r]$ and $\mathcal{F}_{\text{supp}}(\Lambda) \not\sqsubseteq [\pi_1 G_{r-1}]$.*

□

1.6.2 Tiles and applications.

We review from Section 3 of [BFH00] the relations between attracting laminations and tiles of EG strata. In particular Fact 1.56 states the relative train track characterization of attracting laminations. We give some other applications regarding weak attraction, principal rays, and generic leaves.

Fixed notation: For the remainder of the section we fix $\phi \in \text{Out}(F_n)$ acting as the identity on $\mathcal{L}(\phi)$, and we fix $f: G \rightarrow G$ a relative train track representative of ϕ . The conclusions of Fact 1.54 therefore hold, in particular f is EG-aperiodic. We also fix $\Lambda^+ \in \mathcal{L}(\phi)$ and its corresponding EG-stratum $H_r \subset G$.

Recall from Definition 3.1.7 of [BFH00] that for any integer $k \geq 1$, a k -tile of height r is a path of the form $f_\#^k(E)$, where E is an edge of H_r ; we may drop k and/or r if they are clear from the context.

Fact 1.55. *With the notation fixed as above we have:*

- (1) *For each k , each generic leaf of Λ^+ has a decomposition into subpaths each of which is either a k -tile of height r or a path in G_{r-1} .*
- (2) *Let $\{\mu_i\}$ be the collection of maximal subpaths of G_{r-1} that occur in 1-tiles and suppose that τ_0 is a maximal subpath of G_{r-1} in a k -tile τ . Then $\tau_0 = f_\#^m(\mu_i)$ for some $m \leq k-1$ and some μ_i ; moreover, there are m -tiles immediately preceding and following τ_0 in τ .*
- (3) *Each generic leaf of Λ^+ can be exhausted by tiles of height r , meaning that it can be written as an increasing union of subpaths each of which is a tile of height r .*
- (4) *There exists p such that each for each $k \geq i \geq 0$, each $k+p$ -tile of height r contains each i tile of height r .*

Proof. Items (1) and (3) are contained in Lemma 3.1.10 and corollary 3.1.11 of [BFH00]. Item (2) follows from Lemma 5.8 of [BH92] and an obvious induction argument. To prove item (4), choose p so that the p^{th} power of the transition matrix for H_r is positive, and so

for each $i \geq 0$ the $(p+i)^{\text{th}}$ power is also positive; item (4) for $k=0$ follows immediately, and induction establishes it for all higher values of k . \square

The following fact gives the relative train track characterization of the set of attracting laminations of ϕ . In the logical structure of [BFH00] this fact is very closely intertwined with Fact 1.54, but we state it here as an explicit, separate observation in its own right. Recall from Section 1.1.5 the definition of weak attraction for the action of $f_\#$ on $\widehat{\mathcal{B}}(G)$.

Fact 1.56. *Continuing with the notation fixed as above, a line $\ell \in \mathcal{B}$ is a leaf of Λ^+ if and only if some (any) edge of H_r is weakly attracted to ℓ under the action of $f_\#$ on $\widehat{\mathcal{B}}(G)$.*

Proof. For the “only if” direction, consider first a generic leaf ℓ of Λ^+ . By Fact 1.55 (3) any subpath of ℓ is a subpath of some k -tile, which by Fact 1.55 (4) is a subpath of some tile $f_\#^{k+p}(E)$ for some edge E of H_r . It follows that E is weakly attracted to ℓ by iteration of ϕ , and so E is also weakly attracted to every other leaf of Λ^+ , each of which is a weak limit of ℓ .

For the “if” direction, consider any edge E of H_r , and suppose E is weakly attracted to a line ℓ . Choose a generic leaf β of Λ^+ . Each subpath of ℓ is contained in $f_\#^k(E)$ for some k , which by Fact 1.55 (4) is contained in any r -tile of the form $f_\#^{k+p}(E')$. By Fact 1.55 (1), some such tile $f_\#^{k+p}(E')$ is a subpath of β , proving that ℓ is in the weak closure of β which is just Λ^+ . \square

Next we characterize those paths and circuits in G which are weakly attracted to Λ_+ .

Fact 1.57. *Continuing with the notation fixed as above we have:*

- (1) *Given a path or circuit $\sigma \subset G$ the following are equivalent:*
 - (a) *σ is weakly attracted to Λ^+ under the action of f of $\widehat{\mathcal{B}}(G)$.*
 - (b) *There exists $k \geq 0$ and a splitting of $f_\#^k(\sigma)$, one term of which is an edge in H_r .*
 - (c) *For all sufficiently large k , $f_\#^k(\sigma)$ is completely split and one term of the complete splitting is an edge of H_r .*
- (2) *No conjugacy class in F_n is weakly attracted to Λ^+ by iteration of ϕ^{-1} . More generally, letting U be an attracting neighborhood of Λ^+ , for any conjugacy class $[a]$ there are only finitely many values of $k \geq 0$ for which $\phi^{-k}[a] \in U$.*

Proof. The equivalence of (1a) and (1b) is in Corollary 4.2.4 of [BFH00], and the equivalence of (1b) and (1c) follows from Fact 1.34 and the evident fact that for each edge E of H_r and each $i \geq 1$ some term of the complete splitting of $f_\#^i(E)$ is an edge of H_r . To prove item (2), let c be the circuit in G representing $[a]$. Since ϕ fixes Λ^+ , we have $\phi(U) \subset U$. By Lemma 3.1.16 of [BFH00] a generic leaf ℓ is not a circuit, and combined with the fact that $\{\phi^k(U)\}$ is a weak neighborhood basis of ℓ , it follows that there exists I such that $[a] \notin \phi^I(U)$, and so $\phi^{-i}[a] \notin U$ for $i \geq I$. \square

The next fact contains applications of tiles to the understanding of rays in leaves.

Corollary 1.58. *Continuing with the notation fixed above, for any leaf ℓ of Λ^+ , and for any subray $R \subset \ell$ whose initial vertex is principal and whose initial edge is in H_r , R is a principal ray if and only if R is the union of a nested increasing sequence of tiles in ℓ .*

Proof. The “only if” direction follows from the definition of a principal ray. Suppose R is the union of oriented tiles $\tau_1 \subset \tau_2 \subset \tau_3 \subset \dots \subset \ell$ where τ_i is an m_i tile, then $m_i \rightarrow +\infty$ as $i \rightarrow +\infty$. Eventually τ_i contains the initial point of R , and so after throwing away finitely many terms, each τ_i is an initial segment of R . There is an m_0 so that the initial direction of each m_0 -tile is fixed. After passing to a subsequence, there is an m_0 -tile τ_0 such that $\tau_i = f_{\#}^{m_i - m_0}(\tau_0)$ for each $i \geq 1$, and so R is the principal ray associated to the initial direction of τ_0 . \square

The following fact characterizes generic leaves in terms of their asymptotic behavior.

Fact 1.59. *Continuing with the notation fixed above, for any leaf ℓ of Λ^+ , the following are equivalent:*

- (1) ℓ is a generic leaf of Λ^+ .
- (2) ℓ is birecurrent and has height r .
- (3) Both ends of ℓ have height r .

Proof. The implications (1) \implies (2) \implies (3) are obvious, and (2) \implies (1) follows from [BFH00] Lemma 3.1.15.

We prove (3) \implies (2). Consider a generic leaf β of Λ^+ . For each $k \geq 0$, applying Fact 1.55 (1) to β it follows that there exists $N_k > 0$ so that any subpath of β that contains at least N_k edges of H_s contains a k -tile. Since every subpath of ℓ occurs as a subpath of β , every subpath of ℓ that contains at least N_k edges of H_s contains a k -tile. Since both ends of ℓ have height s it follows that every initial or terminal ray of ℓ contains a k -tile for all k . By Fact 1.55 (3), each subpath of β , and hence each subpath σ of ℓ , is contained in a subpath which is a tile, and hence σ occurs as a subpath of every initial or terminal ray of ℓ . This proves that ℓ is birecurrent. \square

1.6.3 EG principal rays and Fix_N

For any outer automorphism ϕ , any expanding lamination, any generic leaf ℓ which ϕ fixes preserving orientation, and any $\Phi \in \text{Aut}(F_n)$ representing ϕ and fixing a lift $\tilde{\ell}$ of ℓ , both endpoints of $\tilde{\ell}$ are contained in $\text{Fix}_+(\widehat{\Phi})$. This fact, reminiscent of Nielsen theory, is stated in [FH11] Remark 3.6. We need a version of this fact which starts not with a generic leaf but only with an EG principal ray R . We will need this fact in the proof of Lemma II.2.11.

For the statement we need a definition. Consider a CT $f: G \rightarrow G$ and a linear NEG edge E with $f(E) = Ew^k$ as in (Linear Edges) in the Definition of a CT ([FH11] Definition 4.7, or Definition 1.29). The ray generated by E has the form $Ew^\infty = Ewww\dots$, called the linear ray generated by E . The ray $Ew^{-\infty} = Ew^{-1}w^{-1}w^{-1}\dots$ will be called the *exceptional ray* of E . Note that the open ray obtained from an exceptional ray by removing its initial point is a nested increasing union of pretrivial paths.

Lemma 1.60. Consider a rotationless $\phi \in \text{Out}(F_n)$, a CT representative $f: G \rightarrow G$, and an EG-stratum H_r with associated $\Lambda \in \mathcal{L}(\phi)$. For any principal direction $E \subset H_r$ generating a principal ray R in G with initial vertex v , there exists a leaf of Λ of the form $\ell = \overline{R}'\gamma R$ such that $f_*(\ell) = \ell$ preserving orientation, γ is a (possibly trivial) Nielsen path in G_{r-1} , and one of the following holds:

- (1) R' is the ray generated by some nonfixed edge $E' \subset G_r$ with fixed initial direction.
- (2) R' is the exceptional ray of some linear edge $E' \subset G_{r-1}$.
- (3) γ is trivial and R' splits into Nielsen paths and fixed edges contained in G_{r-1} .

Furthermore, for any $\Phi \in P(\phi)$ with associated principle lift $\tilde{f}: \widetilde{G} \rightarrow \widetilde{G}$ fixing a lift \tilde{v} of v , letting \widetilde{R} be the lift of R with initial vertex \tilde{v} , letting $\xi \in \partial F_n$ be the limit point of \widetilde{R} , letting $\tilde{\ell}$ be the lift of ℓ containing \widetilde{R} , and letting ξ' be the opposite end of $\tilde{\ell}$, we have $\tilde{f}_*(\tilde{\ell}) = \tilde{\ell}$ preserving orientation, and furthermore:

- (4) $\xi \in \text{Fix}_+(\widehat{\Phi}); \xi' \in \text{Fix}_N(\widehat{\Phi});$ and $\xi' \in \text{Fix}_+(\widehat{\Phi})$ if and only if E' is a principal direction.

Proof. Once (1)–(3) are proved, the rest is proved by writing $\tilde{\ell} = \overline{\widetilde{R}}' \cdot \tilde{\gamma} \cdot \widetilde{R}$ and applying Fact 1.49.

We turn now to the construction of the leaf ℓ . Define a *split leaf segment of Λ in G* to be a completely split path δ in G_r such that some generic leaf ℓ of Λ splits as $\ell = \bar{\rho}_- \cdot \delta \cdot \rho_+$. We use without comment that the f_* image of a split leaf segment is a split leaf segment, and that every concatenation of consecutive terms of the complete splitting of a split leaf segment, which we shall call “a split leaf subsegment”, is a split leaf segment. Since each generic leaf of Λ crosses every edge of H_r infinitely often, there exists a split leaf segment of the form $\overline{E}'' \cdot \alpha \cdot E$ such that E'' is an edge of H_r and α is a (possibly trivial) completely split path in G_{r-1} .

Suppose first that α is trivial or a Nielsen path. After replacing $\overline{E}'' \cdot \alpha \cdot E$ by its image under some iterate \tilde{f}_k^* so that $\tilde{f}_k^*(E'')$ has fixed initial direction, and then passing to a split leaf subsegment, we may assume that E'' has fixed initial direction and so is a principal direction. Letting $E' = E''$ and letting R' be the ray generated by E' , the line $\ell = \overline{R}'\alpha R$ satisfies (1) with $E' \subset H_r$.

Suppose now that α is nontrivial and not a Nielsen path. Discarding E'' we narrow our focus to split leaf segments of the form $\alpha \cdot E$ where α is a completely split path in G_{r-1} and is not a Nielsen path. There is a unique splitting $\alpha = \beta_2 \cdot \beta_1 \cdot \gamma$ such that γ is the longest Nielsen path which is a terminal segment of the complete splitting of α , and β_1 is the term of the complete splitting preceding γ ; note that β_2 may be trivial but β_1 is nontrivial. Amongst all such paths $\alpha \cdot E = \beta_2 \cdot \beta_1 \cdot \gamma \cdot E$ under consideration, let $s < r$ be the minimum height of β_1 . Discarding β_2 we further narrow our focus to split leaf segments of the form $\beta_1 \cdot \gamma \cdot E$ where γ is a Nielsen path in G_{r-1} and β_1 is a height s complete splitting term which is neither a fixed edge nor an indivisible Nielsen path. Note also that β_1 is not contained in a zero stratum, because by applying (Zero Strata) of the definition of a CT ([FH11] Definition 4.7, or Definition 1.29) it follows that no vertex in a zero stratum is fixed by f , but the terminal point of β_1 , which equals the initial point of γ , is fixed. We now go through the remaining cases of β_1 .

If $\beta_1 = \overline{E}'$ where E' is an EG or NEG edge with fixed initial direction then, taking R' to be the ray generated by E' , we are done with $\ell = \overline{R}' \cdot \gamma \cdot R$ satisfying (1). If $\beta_1 = E'$ where E' is an EG edge with nonfixed terminal direction then, after replacing $\beta_1 \cdot \gamma \cdot E$ by its image under some iterate $\tilde{f}_\#^k$ so that $\tilde{f}_\#^k(E')$ has fixed terminal direction, and then passing to a split leaf subsegment, we reduce to the previous case. If $\beta_1 = E_s$ is a nonfixed, nonlinear NEG edge with fixed initial direction then $f(\beta_1) = E_s \cdot u$ where u is of height $< s$ and is not a Nielsen path, and so after one replacing $\beta_1 \cdot \gamma \cdot E$ with its image under $\tilde{f}_\#^k$ and passing to a split leaf subsegment, we obtain an example with smaller s , a contradiction. If $\beta_1 = E_s$ is a nonfixed, linear NEG edge with fixed initial direction and $f(E_s) = E_s \cdot u$ then $f_\#^k(\beta_1 \cdot \gamma \cdot E)$ contains a split leaf subsegment of the form $u^k \cdot \gamma \cdot f_\#^k(E)$ and passing to a weak limit one obtains $\ell = \overline{R}' R$ satisfying (3) with $R' = \bar{\gamma}u^{-\infty}$. In the remaining case wheree $\beta_1 = E_i w^p \overline{E}_j$ is an exceptional path, with $f(E_i) = E_i w^{d_i}$, $f(E_j) = E_j w^{d_j}$, then $f_\#^k(\beta_1 \cdot \gamma \cdot E)$ contains a split leaf subsegment of the form $E_i w^{p+k(d_i-d_j)} \overline{E}_j \cdot \gamma \cdot f_\#^k(E)$, and passing to a weak limit one obtains $\ell = \overline{R}' \cdot \gamma \cdot R$ where R' is one of the two rays $E_j w^{-\infty}$ or $E_j w^{+\infty}$ depending on whether $d_i - d_j$ is positive or negative; this is a linear ray if $d_j > 0$ and an exceptional ray otherwise. \square

1.6.4 Pushing forward attracting laminations.

One commonly used result about attracting laminations says that for each $\Phi \in \text{Aut}(F_n)$ and each Φ -invariant free factor of A , the inclusion $A \hookrightarrow F_n$ induces an injection from the set of attracting laminations of the outer automorphism class of $\Phi \mid A \in \text{Aut}(A)$ to the set of attracting laminations for the outer automorphism class of $\Phi \in \text{Aut}(F_n)$ itself. This is proved by an easy relative train track argument, using that every outer automorphism preserving a free factor is represented by a relative train track map having a filtration element representing that free factor. We need a generalization of this result which applies without the assumption that A is a free factor; in the conclusion we must sacrifice injectivity of the induced map.

Given a finite rank subgroup $A < F_n$, the embedding $\partial A \hookrightarrow \partial F_n$ induces an embedding $\tilde{\mathcal{B}}(A) \hookrightarrow \tilde{\mathcal{B}}(F_n)$ which in turn induces a continuous map $\beta_A: \mathcal{B}(A) \rightarrow \mathcal{B}(F_n)$, defined by saying that the β_A -image of an A -orbit in $\tilde{\mathcal{B}}(A)$ is the unique F_n -orbit in $\tilde{\mathcal{B}}(F_n)$ containing the given A -orbit.

The following lemma will be used in the proof of Proposition 2.11.

Lemma 1.61. *For each $\phi \in \text{Out}(F_n)$, each representative $\Phi \in \text{Aut}(F_n)$, and each finite rank Φ -invariant subgroup $A < F_n$, letting $\phi \mid A \in \text{Out}(A)$ denote the outer automorphism class of $\Phi \mid A$, for each dual lamination pair $\Lambda^\pm \in \mathcal{L}^\pm(\phi \mid A)$ the following hold:*

- (1) $\beta_A(\Lambda^+) \in \mathcal{L}(\phi)$ and $\beta_A(\Lambda^-) \in \mathcal{L}(\phi^{-1})$
- (2) *Letting $B < A$ represent the free factor support of Λ^\pm in A , the free factor supports of $\beta_A(\Lambda^+)$ and of $\beta_A(\Lambda^-)$ in F_n are each equal to the free factor support of B in F_n .*
- (3) $\beta_A(\Lambda^+), \beta_A(\Lambda^-)$ are a dual lamination pair in $\mathcal{L}^\pm(\phi)$.

Note that unlike Fact 1.4 the notation $\phi \mid A$ does not presuppose well-definedness independent of the choice of Φ . However, it is at least well-defined up to post-composing Φ by an inner automorphism determined by an element of A .

Proof. Since duality of laminations is defined by having the same free factor support, clearly (2) \implies (3).

Also (1) \implies (2), which we see as follows. We add a superscript “ A ” to certain notations, e.g. $[\cdot]^A$ denotes conjugacy classes in A and $\mathcal{F}_{\text{supp}}^A(\cdot)$ denotes free factor support in A ; with no superscript these notations retain their ordinary meaning in F_n . By Fact 1.11, $\mathcal{F}_{\text{supp}}(B)$ is the free factor support of the set of lines carried by the one component subgroup system $[B]$, and so $\beta_A(\Lambda^+)$ is carried by $\mathcal{F}_{\text{supp}}(B)$, proving that $\mathcal{F}_{\text{supp}}(\beta_A(\Lambda^+)) \sqsubset \mathcal{F}_{\text{supp}}(B)$. For the opposite direction, suppose that $\mathcal{C} = \{[C_1], \dots, [C_K]\}$ is a free factor system in F_n that supports $\beta_A(\Lambda^+)$. By the Kurosh Subgroup Theorem the collection of conjugacy classes $\{[A \cap C_k^g]^A \mid g \in F_n, 1 \leq k \leq K\}$ forms a free factor system in A , and by supposition it carries Λ^+ in A . It follows that $\{[B]^A\} = \mathcal{F}_{\text{supp}}^A(\Lambda^+) \sqsubset \mathcal{C}$ which implies that $B \subset C_k^g$ for some $k = 1, \dots, K$ and some $g \in F_n$, which implies that $\mathcal{F}_{\text{supp}}(B) \sqsubset \mathcal{C}$. Since this is true for all \mathcal{C} supporting $\beta_A(\Lambda^+)$ we have $\mathcal{F}_{\text{supp}}(B) \sqsubset \mathcal{F}_{\text{supp}}(\beta_A(\Lambda^+))$, and so $\mathcal{F}_{\text{supp}}(\beta_A(\Lambda^+)) = \mathcal{F}_{\text{supp}}(B)$. The same is true for Λ^- by replacing ϕ with ϕ^{-1} , which proves (2).

We prove (1) just for Λ^+ ; the same follows for Λ^- by replacing ϕ with ϕ^{-1} . For the proof we drop the “ $+$ ” superscript and simply write Λ for Λ^+ . We adapt the proof of [BFH00] Lemma 3.1.9. We may freely pass to any positive power of ϕ , and so in particular we may assume $\phi \mid A$ fixes Λ and so ϕ fixes $\beta_A(\Lambda)$.

Pick a relative train track map $h: H \rightarrow H$ representing $\phi \mid A$ with filtration elements H_i and strata $H_i \setminus H_{i-1}$, having in particular an aperiodic EG stratum $H_r \setminus H_{r-1}$ corresponding to Λ . After passing to a positive power, there is an edge $E \subset H_r \setminus H_{r-1}$ and a point $x \in \text{int}(E)$ such that $h(x) = x$. Up in the universal cover choose lifts $\tilde{E} \subset H$ of E , $\tilde{x} \in \text{int}(\tilde{E})$ of x , and $\tilde{h}: \tilde{H} \rightarrow \tilde{H}$ of h so that $\tilde{h}(\tilde{x}) = \tilde{x}$. After post-composing Φ by an inner automorphism determined by an element of A , we may assume that \tilde{h} is the lift of h corresponding to $\Phi \mid A$. Letting $\tilde{\ell} = \cup_{i=1}^{\infty} \tilde{h}_\#^k(E) \in \tilde{\mathcal{B}}(H)$, its downstairs image $\ell \in \mathcal{B}(H)$ is a generic leaf of the realization of Λ in H , as shown in the proof of [BFH00] Lemma 3.1.9. By definition the leaf $\ell \in \Lambda$ is birecurrent in $\mathcal{B}(A)$ and dense in Λ , and these evidently imply that the leaf $\beta_A(\ell) \in \beta_A(\Lambda)$ is birecurrent in $\mathcal{B}(F_n)$ and dense in $\beta_A(\Lambda)$. It remains to prove existence of an attracting neighborhood of $\beta_A(\ell)$ in $\mathcal{B}(F_n)$ under the action of ϕ .

The leaf $\tilde{\ell}$ crosses edges of $\tilde{H}_r \setminus \tilde{H}_{r-1}$ infinitely often on both directions; enumerate those crossings in order as \tilde{E}_k , $k \in \mathbb{Z}$, with $\tilde{E} = \tilde{E}_0$. For each integer $k \geq 0$ let $\tilde{\ell}_k$ be the subsegment of $\tilde{\ell}$ starting with \tilde{E}_{-k} and ending with \tilde{E}_k , and so $\cup_k \tilde{\ell}_k = \tilde{\ell}$. Let $\ell_k \in \widehat{\mathcal{B}}(H)$ be the downstairs image of $\tilde{\ell}_k$, so ℓ is a weak limit of ℓ_k as $k \rightarrow +\infty$. After passing to a positive power we may assume that the $h_\#$ -image of each edge of $H_r \setminus H_{r-1}$ crosses at least two edges, and we obtain the inclusion $\tilde{h}_\#(\tilde{\ell}_k) \supset \tilde{\ell}_{2k}$ (c.f. [BFH00] Lemma 2.5.1). From this inclusion, together with Lemma 1.6 (1), it follows that

(*) For each $d \geq 0$, if $k \geq \text{BCC}(h) + d$ then $\tilde{h}_\#(\tilde{\ell}_k) \supset \tilde{\ell}_{k+d}$.

because both components of $\tilde{\ell}_{2k} - \tilde{\ell}_{k+d}$ have length at least $k - d \geq \text{BCC}(h)$.

Property (*) with $d = 1$ is a key step in the proof of [BFH00] Lemma 3.1.9 for showing that ℓ has an attracting neighborhood under iteration of $\phi \mid A$. The strategy of our proof that $\beta_A(\ell)$ has an attracting neighborhood under iteration of ϕ is to use bounded cancellation arguments to transfer this key step over to the setting of a topological representative of ϕ .

Pick an F_n -marked graph G and a π_1 -injective map $\mu: H \rightarrow G$ which, with appropriate choices of bases points and paths between them, induces the injection $A \hookrightarrow F_n$. Let $\tilde{\mu}: \tilde{H} \rightarrow \tilde{G}$ be an A -equivariant lift of μ . Pick a homotopy equivalence $g: G \rightarrow G$ representing ϕ , and so $g \circ \mu$ is homotopic to $\mu \circ h$. Let $\tilde{g}: \tilde{G} \rightarrow \tilde{G}$ be the lift corresponding to the automorphism Φ . It follows that the homotopy between the maps $g \circ \mu$ and $\mu \circ h: H \rightarrow G$ lifts to an A -equivariant homotopy between the maps $\tilde{g} \circ \tilde{\mu}$ and $\tilde{\mu} \circ \tilde{h}: \tilde{H} \rightarrow \tilde{G}$. By compactness of H , the track of each point under this homotopy has diameter in \tilde{G} bounded by a uniform constant T .

Working with the realizations of Λ in H and of $\beta_A(\Lambda)$ in G , the realizations in G of the leaves of $\beta_A(\Lambda)$ are precisely the $\mu_{\#}$ -images of the realizations in H of the leaves of Λ . In particular, $\ell' = \mu_{\#}(\ell)$ is the realization in G of $\beta_A(\ell) \in \beta_A(\Lambda)$. We must prove that ℓ' has an attracting neighborhood in $\mathcal{B}(G)$ under the action of $g_{\#}$.

In \tilde{G} let $\tilde{\ell}'_k = \tilde{\mu}_{\# \#}(\tilde{\ell}_k)$ and in G let $\ell'_k = \mu_{\# \#}(\ell_k)$. Applying Lemma 1.6 (3) we have a nested sequence $\tilde{\ell}'_1 \subset \tilde{\ell}'_2 \subset \dots$ whose union is contained in $\tilde{\ell}' = \tilde{\mu}_{\#}(\tilde{\ell})$. Letting $V'_k = V(\ell'_k; G)$ we obtain a nested sequence of weakly open sets $V'_1 \supset V'_2 \supset \dots$ of $\mathcal{B}(G)$. We shall prove:

(a) The sequence $V'_1 \supset V'_2 \supset \dots$ is a weak neighborhood basis for ℓ' .

(b) For all sufficiently large k , say $k \geq K$, we have $g_{\#}(V'_k) \subset V'_{k+1}$.

from which it immediately follows that V'_K is a weakly attracting neighborhood of ℓ' under iteration of $g_{\#}$.

We claim that:

(**) For each $M \geq 0$ there exists e_M such that if $e \geq e_M$ then for each k both components of $\tilde{\ell}'_{k+e} - \tilde{\ell}'_k$ have length $> M$.

To prove the claim, from the nesting of these arcs we obtain a bijection correspondence between the endpoints of $\tilde{\ell}'_{k+e}$ and the endpoints of $\tilde{\ell}'_k$, so it is sufficient to prove that the distance between corresponding endpoints is $> M$. We have $\tilde{\ell}'_{k+e} - \tilde{\ell}'_k = \tilde{\mu}_{\# \#}(\tilde{\ell}_{k+e}) - \tilde{\mu}_{\# \#}(\tilde{\ell}_k)$. For each arc $\alpha \subset \tilde{H}$, by Lemma 1.6 (1) the arc $\tilde{\mu}_{\# \#}(\alpha) \subset \tilde{G}$ is obtained from the arc $\tilde{\mu}_{\#}(\alpha)$ by trimming away at most $\text{BCC}(\mu)$ from both ends of $\tilde{\mu}_{\#}(\alpha)$, and so we get a bijective correspondence between the endpoints of $\tilde{\mu}_{\#}(\alpha)$ and those of $\tilde{\mu}_{\# \#}(\alpha)$ so that corresponding endpoints have distance $\leq \text{BCC}(\mu)$. It therefore suffices to show that corresponding endpoints of the arcs $\tilde{\mu}_{\#}(\tilde{\ell}_{k+e})$, $\tilde{\mu}_{\#}(\tilde{\ell}_k)$ have distance $> M - 2\text{BCC}(\mu)$. Using that $\tilde{\mu}: \tilde{H} \rightarrow \tilde{G}$ is an (a, c) quasi-isometric embedding for some $a \geq 1$, $c \geq 0$, it is sufficient to show that corresponding endpoints of $\tilde{\ell}_{k+e}$, $\tilde{\ell}_k$ have distance $> a(M - 2\text{BCC}(\mu) + c)$, which is guaranteed by choosing $e > e_M = a(M - 2\text{BCC}(\mu) + c)$.

From the claim it follows that $\cup_k \tilde{\ell}'_k = \ell'$, and so $V'_1 \supset V'_2 \supset \dots$ is a neighborhood basis of ℓ' in $\mathcal{B}(G)$, verifying (a). We shall also use the claim, with an appropriate choice of M , to verify (b).

We have:

$$\tilde{\ell}'_{k+d} = \tilde{\mu}_{\# \#}(\tilde{\ell}_{k+d}) \subset \tilde{\mu}_{\# \#}(h_{\# \#}(\tilde{\ell}_k))$$

which follows by applying (*) and Lemma 1.6, as long as $k \geq K_d = \text{BCC}(h) + d$. We make an appropriate choice of d below. Applying Lemma 1.6 again we have:

$$\tilde{\ell}'_{k+d} \subset (\tilde{\mu} \circ \tilde{h})_{\# \#}(\tilde{\ell}_k) \subset (\tilde{\mu} \circ \tilde{h})_{\#}(\tilde{\ell}_k)$$

Let $M = \max\{T, \text{BCC}(g \circ \mu)\}$, choose $e \geq e_M$ and let $d = 2e + 1$, and so

$$\tilde{\ell}'_{k+2e+1} \subset (\tilde{\mu} \circ \tilde{h})_*(\tilde{\ell}_k)$$

Applying $(**)$ twice, using $M \geq T$ the first time and $M \geq \text{BCC}(g \circ \mu)$ the second, together with a few more applications of Lemma 1.6, we have:

$$\begin{aligned}\tilde{\ell}'_{k+e+1} &\subset (\tilde{g} \circ \tilde{\mu})_*(\tilde{\ell}_k) \\ \tilde{\ell}'_{k+1} &\subset (\tilde{g} \circ \tilde{\mu})_{\#}(\tilde{\ell}_k) \subset \tilde{g}_{\#}(\tilde{\mu}_{\#}(\tilde{\ell}_k)) = \tilde{g}_{\#}(\tilde{\ell}'_k)\end{aligned}$$

and so $g_{\#}(V'_k) \subset V'_{k+1}$. \square

2 Geometric EG strata and geometric laminations

A geometric outer automorphism $\phi \in F_n$, as defined in Section 4 of [BH92], is one modeled by a pseudo-Anosov homeomorphism $f: S \rightarrow S$ of a surface with nonempty boundary. The broader concept of a geometric EG stratum $H_r = G_r \setminus G_{r-1}$ of a relative train track map $f: G \rightarrow G$ was introduced in Definition 5.1.4 of [BFH00]: roughly speaking H_r is geometric if the restriction $f: G_r \rightarrow G_r$ is modeled by a 2-dimensional dynamical system obtained by gluing together the restriction $f: G_{r-1} \rightarrow G_{r-1}$ with a pseudo-Anosov surface homeomorphism $h: S \rightarrow S$, by attaching all but one component of ∂S to G_{r-1} .

In Section 2.1 we review and reformulate the definition of geometric strata from Definition 5.1.4 of [BFH00], couching it in terms of existence of the 2-dimensional dynamical system alluded to above, which is formally described in Definitions 2.1 and 2.2 where it is dubbed a *geometric model* for f and H_r .

In Section 2.2 we study a natural graph of groups decomposition of F_n associated to a geometric stratum called the *peripheral splitting*, which records group theoretic information underlying the topological manner in which the geometric model is glued together from the surface S and the graph G .

In Sections 2.3 and 2.4 we study the dual lamination pair of a geometric stratum. Proposition 2.11, which is based on the geometric case of the proof of Proposition 6.0.8 of [BFH00], shows that the laminations of a geometric stratum H_r can be identified with the stable and unstable laminations of the pseudo-Anosov homeomorphism $h: S \rightarrow S$. Lemma 2.12 shows that the duality relation between attraction laminations of an outer automorphism and those of its inverse preserves inclusion of laminations. Proposition 2.13 is a new result which proves that “geometricity” is actually an invariant of an attracting lamination, indeed of a dual lamination pair, not just of a stratum.

In Sections 2.5 and 2.6 we establish some complementarity properties for geometric strata that are analogous to properties of subsurfaces and their complements. The results of these sections will be used in Part II [HM13b] to establish certain invariance properties of elements of $\text{IA}_n(\mathbb{Z}/3)$.

2.1 Defining and characterizing geometric strata

In this section we review the definition of geometric strata, and other results from [BFH00] concerning their properties. In particular, given a CT $f: G \rightarrow G$ and an EG stratum H_r , geometricity of H_r is defined in Definition 5.1.4 of [BFH00] in terms of an object that we term a “weak geometric model”, formalized in Definition 2.1. In Definition 2.4 we reformulate geometricity again in terms of a (stronger) “geometric model”. Fact 2.3 asserts the equivalence of these and other formulations of geometricity of an EG stratum. The proof of Fact 2.3, which follows several results of [BFH00], takes up subsections 2.1.3 and 2.1.4.

2.1.1 Defining weak geometric models and geometric strata.

The definitions of various 1-dimensional topological representatives of elements of $\text{Out}(F_n)$ — train track maps and relative train track maps [BH92], improved relative train track maps [BFH00], and CTs [FH11]— can each be broken into two parts: static data consisting of a filtered marked graph satisfying certain conditions; and dynamic data consisting of a homotopy equivalence of that graph, again satisfying various conditions.

Given a CT $f: G \rightarrow G$ and an EG stratum $H_r \subset G$, Definition 2.1 of a weak geometric model and 2.4 of a geometric model, are similarly broken up into static data and dynamic data. The static data describes how to glue up a 2-complex from a surface and a graph, and how to mark that 2-complex by a certain homotopy equivalence to G . The dynamic data describes a dynamical system on the 2-complex obtained by gluing together a homeomorphism of the surface and a homotopy equivalence of the graph, with conditions that describe relations amongst the parts of the dynamical system and the given CT f , these relations all expressed as commutative or homotopy commutative diagrams.

Notation fixed for the remainder of Section 2.1. We fix the following:

- A rotationless $\phi \in \text{Out}(F_n)$;
- A CT $f: G \rightarrow G$ representing $\phi \in \text{Out}(F_n)$;
- An EG stratum $H_r \subset G$.

Definition 2.1. Weak geometric model for the EG stratum H_r of $f: G_r \rightarrow G_r$.

The static data of the weak geometric model is as follows:

- (1) We are given a compact, connected surface S of negative Euler characteristic having nonempty boundary with components $\partial S = \partial_0 S \cup \dots \cup \partial_m S$ ($m \geq 0$). We refer to $\partial_0 S$ as the *upper boundary* of S and to $\partial_1 S, \dots, \partial_m S$ as the *lower boundaries*.
- (2) For each lower boundary $\partial_i S$, $i = 1, \dots, m$, we are given a homotopically nontrivial closed edge path $\alpha_i: \partial_i S \rightarrow G_{r-1}$. The map α_i need not be a local embedding.
- (3) We let Y be the 2-complex defined as the quotient of the disjoint union $S \amalg G_{r-1}$ obtained by gluing each lower boundary circle $\partial_i S$ to G_{r-1} using α_i as a gluing map. We let $j: S \amalg G_{r-1} \rightarrow Y$ denote the quotient map.

Several subsets of the domain of the quotient map $j: S \amalg G_{r-1} \rightarrow Y$ are topologically embedded in Y , and we will by convention identify these subsets with their images in Y .

These include: G_{r-1} and $\partial_0 S$, each identified with a subcomplex of Y ; the noncompact surface $\text{int}(S)$ is identified with an open subset of Y ; and the noncompact surface-with-boundary $\text{int}(S) \cup \partial_0 S$ is identified with an open subset of Y .

- (4) We are given an embedding $G_r \hookrightarrow Y$ which extends the embedding $G_{r-1} \hookrightarrow Y$, and we identify G_r with its image in Y . We are also given a deformation retraction $d: Y \rightarrow G_r$. These satisfy the following properties:
 - (a) $G_r \cap \partial_0 S$ is a single point of Y denoted p_r , and the restriction $d|_{\partial_0 S}$ with base point p_r is a closed, indivisible Nielsen path of f of height r denoted ρ_r .
 - (b) The composition $G_{r-1} \xrightarrow{j} Y \xrightarrow{d} G_r$ is the inclusion map.
 - (c) The interior of H_r in G_r equals $H_r - G_{r-1} = H_r \cap (Y - G_{r-1}) = H_r \cap (\text{int}(S) \cup \partial_0 S)$.

Items (4b), (4c) follow immediately from the previous items; we include them here for clarity at the risk of redundancy.

The dynamic data of the weak geometric model is as follows:

- (5) We are given a homotopy equivalence $h: Y \rightarrow Y$, and we are given a homeomorphism $\Psi: (S, \partial_0 S) \rightarrow (S, \partial_0 S)$ with pseudo-Anosov mapping class $\psi \in \mathcal{MCG}(S)$, subject to the following compatibility conditions:
 - (a) The maps $(f|_{G_r}) \circ d$ and $d \circ h: Y \rightarrow G_r$ are homotopic.
 - (b) The maps $j \circ \Psi$ and $h \circ j: S \rightarrow Y$ are homotopic.

To summarize, a weak geometric model of the CT $f: G_r \rightarrow G_r$ for the EG stratum H_r is a tuple of data $(Y, d, S, (\partial_i S)_{i=0}^m, j, (\alpha_i)_{i=1}^m, h, \Psi)$ as described and satisfying the conditions above. Our usual terminology will suppress all of the data except Y , saying that Y is a *weak geometric model of $f: G_r \rightarrow G_r$ for H_r* (or “with respect to H_r ”, or other such phrases). When the CT f is understood we simply say Y is a *weak geometric model for H_r* . This completes Definition 2.1.

Definition 2.2. Geometricity of the EG stratum H_r . We say that the EG stratum H_r of the CT $f: G \rightarrow G$ is *geometric* if a weak geometric model of $f: G_r \rightarrow G_r$ for H_r exists.

Remark: Comparing definitions of geometricity: Definition 2.2 versus Definition 5.1.4 of [BFH00]. These two definitions of geometricity of H_r are logically equivalent. The definitions are similarly structured, defining geometricity in terms of the existence of a certain homotopy model for $f: G_r \rightarrow G_r$ which incorporates a pseudo-Anosov surface homeomorphism, but the models used in the two definitions have some differences of expression and mathematical detail. Nonetheless existence of these two models is equivalent, because each is equivalent to existence of a closed, indivisible Nielsen path ρ_r of height r . For Definition 2.1 this is part of Fact 2.3 stated below. For [BFH00] Definition 5.1.4 this equivalence is proved by quoting results of [BFH00] as follows: combining Definition 5.1.4 with Theorem 5.1.5 (eg-iii) one obtains existence of ρ_r ; and in the other direction, if a closed, indivisible Nielsen path ρ_r of height r exists then by Lemma 5.1.7 (or

see Fact 1.42 (2) above) the path ρ_r crosses each edge of H_r exactly twice, and then by Proposition 5.3.1, H_r satisfies Definition 5.1.4.

The reader may also be interested in a more direct comparison between the two types of geometric models. Existence of a geometric model as in Definition 2.1 fairly directly implies existence of the model built into the [BFH00] definition. First, in Definition 2.1 we do not express annulus neighborhoods of $\partial_i S$ as they are expressed in [BFH00], although the structure of those neighborhoods and the requirements thereon are easily recovered. Also, our item (2) does not require the attaching maps α_i to be local embeddings, as they are required to be in [BFH00]; that requirement is easily restored by homotoping the attaching maps α_i , however doing so may destroy the possibility of the embedding $G_r \hookrightarrow Y$. Fortunately, while our item (4) does require this embedding and the deformation retraction $d: Y \rightarrow G_r$, all that is required in place of d in the [BFH00] definition is a homotopy equivalence $Y \xrightarrow{\sim} G_r$. Also, our (5b) requires commutativity only up to homotopy, not commutativity on the nose as required in [BFH00], but such stronger commutativity is also easily recovered. From this discussion it is easy to fill in the details and prove that geometricity of H_r as in Definition 2.2 implies geometricity of H_r as in [BFH00].

In the other direction, most items of Definition 2.2 are recovered from Definition 5.1.4 of [BFH00] in a similarly easy fashion, with the exception of item (4), where we require the embedding $G_r \hookrightarrow Y$ and deformation retraction $d: Y \rightarrow G_r$. The reader who knows the proof of [BFH00] Proposition 5.3.1, or who follows closely the proof of Fact 2.3 which is modeled on [BFH00] Proposition 5.3.1, will see that if one weakens (4) appropriately by replacing the deformation retraction $d: Y \rightarrow G_r$ with a homotopy equivalence as in [BFH00], then the full strength of (4) can nevertheless be recovered (see the “Remark” shortly into the proof of Fact 2.3).

In fact the gist of our proof of Fact 2.3 will be to go through the proof of [BFH00] Proposition 5.3.1, tweaking it and using various of its parts, in order to directly verify item (4) and other items of Definition 2.1.

Fact 2.3 (Characterization of geometric strata). *The following are equivalent:*

- (1) H_r is a geometric stratum.
- (2) There exists a closed, height r indivisible Nielsen path ρ_r .
- (3) There exists a height r indivisible Nielsen path ρ_r which crosses each edge of H_r exactly twice.

Furthermore, if these hold then each component of G_{r-1} is noncontractible.

The proof is found in Section 2.1.4, with preliminary material in Section 2.1.3.

2.1.2 Defining geometric models.

We continue with the notation fixed in Section 2.1.1.

A “weak” geometric model for H_r extends by a unique construction to a (“strong”) geometric model for H_r . The former consists of a certain 2-complex Y and homotopy equivalence $h: Y \rightarrow Y$ which is a homotopy model of the restricted map $f: G_r \rightarrow G_r$, the latter is an extension of $h: Y \rightarrow Y$ to a homotopy model $h: X \rightarrow X$ for the entire map $f: G \rightarrow G$.

Definition 2.4. Geometric model for the EG stratum H_r of $f: G \rightarrow G$.

Given a weak geometric model Y of $f: G_r \rightarrow G_r$ relative to H_r , and incorporating all the notation of Definition 2.1, a geometric model of $f: G \rightarrow G$ relative to H_r is defined as follows:

- (1) Let X be the 2-complex obtained as the quotient of the disjoint union $G \amalg Y$ obtained by identifying the two copies of G_r , one embedded in G and the other embedded in Y , via the identity map on G_r . We identify G, Y with their image subcomplexes in X under the quotient map $G \amalg Y \rightarrow X$.
- (2) Let $d: X \rightarrow G$ be the deformation retraction obtained by extension of the deformation retraction $d: Y \rightarrow G_r$, requiring that the restriction of d to $G \setminus G_r = X \setminus Y$ be the identity.
- (3) Let $h: X \rightarrow X$ be the homotopy equivalence obtained by extension of the homotopy equivalence $h: Y \rightarrow Y$, requiring that $h|_{(G \setminus G_r)} = f|_{(G \setminus G_r)}$. Note that the maps $d \circ h$ and $f \circ d: X \rightarrow G$ are homotopic.

This completes the defining items (1)–(3) of the geometric model. Again, the terminology in full is that “ X (with accompanying data) is a geometric model of $f: G \rightarrow G$ for H_r ”, and in brief we say that “ X is a geometric model for H_r ”. We will also say that $Y \subset X$ is the weak geometric submodel.

We continue with an additional item that follows immediately from the above defining items, plus some definitions/terminologies/notations which are used in connection with geometric models.

- (4) The frontier in X of the subset $\text{int}(S) \subset Y \subset X$ is the disjoint union of $j(\partial S)$ with the finite set

$$\text{int}(S) \cap (\cup_{s>r} H_s) = (H_r \cap (\cup_{s>r} H_s)) - \{p_r\}$$

Each element of the latter set is a vertex of H_r called an *attaching point* of X .

The composition $S \xrightarrow{j} Y \xrightarrow{d} G_r \subset G$ is equal to the composition $S \xrightarrow{j} Y \subset X \xrightarrow{d} G$, and by abusing notation we shall write this map as $d|_S: S \rightarrow G$. Subject to choices of base points and paths between them, the map $d|_S$ induces a homomorphism $d_*: \pi_1 S \rightarrow \pi_1 G \approx F_n$, and an application of Van Kampen’s theorem implies that this homomorphism is injective (or see Lemma 2.7). The image subgroup $d_*(\pi_1 S) < F_n$ has conjugacy class denoted $[\pi_1 S]$ or just $[S]$, and this conjugacy class is well-defined independent of choices of base points and paths. Also, for $0 \leq i \leq m$ the closed curve $d|_{\partial_i S} = \gamma_i$ with its two orientations determines an unordered pair of inverse conjugacy classes in F_n denoted $[\partial_i S]^\pm = \{[\partial_i S], [\partial_i S]^{-1}\}$; the assignment of one as the positive and the other as the negative orientation is usually not relevant, except for situations where an orientation of S and the induced boundary orientation on ∂S is under consideration (see also [FH11] Definition 4.1, “unoriented conjugacy classes”). Note that $[\partial_i S]$ need not be a root-free conjugacy class, and if $i \neq j$ then the conjugacy classes $[\partial_i S]$, $[\partial_j S]$ might have a common power or might even be equal. However, $[\partial_0 S]$ is always root-free and never has a common power with $[\partial_i S]$, $i = 1, \dots, m$. We refer to the set $[\partial S]^\pm = \cup_{i=0}^m [\partial_i S]^\pm$ as the *peripheral conjugacy classes* of

the geometric model. We also say that $[\partial_0 S]^\pm$ are the *top* peripheral conjugacy classes, and that $\cup_{i=1}^m [\partial_i S]^\pm$ are the *bottom* peripheral conjugacy classes. When orientation is irrelevant we shall abuse notation and work with $[\partial S] = \{[\partial_0 S], \dots, [\partial_m S]\}$ instead of with $[\partial S]^\pm$.

This completes Definition 2.4 with all its associated terminology/notation.

2.1.3 Invariant free factor systems associated to a geometric model.

In this section we continue with the notation fixed in Section 2.1.1.

Before starting the proof of Fact 2.3 we state and prove a lemma describing some properties of various ϕ -invariant free factor systems associated to a geometric model, in particular the free factor supports of the peripheral conjugacy classes ∂S and of the subgroup $\pi_1 S$ itself.

For the reader who wishes to get quickly to the proof of Fact 2.3 in Section 2.1.4, only item (2) of Fact 2.5 is needed for that proof. Item (1) will be needed in Section 2.3.2, and item (4) in Part IV [HM13d]. Item (3) will be used without comment in many contexts employing geometric models.

Lemma 2.5. *If X is a geometric model of $f: G \rightarrow G$ and H_r then, adopting all the notation of 2.4, we have:*

$$(1) \quad \mathcal{F}_{\text{supp}}[\partial S] = \mathcal{F}_{\text{supp}}[\pi_1 S]$$

$$(2) \quad \mathcal{F}_{\text{supp}}[\partial_0 S] \not\subset [\pi_1 G_{r-1}]$$

(3) *Each of the following free factor systems is ϕ -invariant:*

$$\mathcal{F}_{\text{supp}}[\pi_1 S] = \mathcal{F}_{\text{supp}}[\partial S], \quad \mathcal{F}_{\text{supp}}[\partial_0 S], \quad \text{and} \quad \mathcal{F}_{\text{supp}}\{[\partial_1 S], \dots, [\partial_m S]\}$$

(4) *There does not exist a ϕ -periodic free factor system properly contained between $\mathcal{F}_{\text{supp}}\{[\partial_1 S], \dots, [\partial_m S]\}$ and $\mathcal{F}_{\text{supp}}[\pi_1 S]$.*

Proof. Let G_r^0 be the component of G_r containing H_r , let Y^0 be the component of Y containing the connected set $j(S)$, and note that the deformation retraction $d: Y \rightarrow G_r$ restricts to a deformation retraction $d': Y^0 \rightarrow G_r^0$. Noting that $F' = \pi_1(Y^0) = \pi_1(G_r^0)$ is a free factor of F_n , clearly we have

$$\mathcal{F}_{\text{supp}}\{[\partial_1 S], \dots, [\partial_m S]\} \subset \mathcal{F}_{\text{supp}}[\partial S] \subset \mathcal{F}_{\text{supp}}(\pi_1 S) \subset [F']$$

For the proofs of (1) and (2), by restricting to F' we may assume that $F' = F_n$ and so $Y = Y^0 = X$ and $G = G_r = G_r^0$.

Item (1) will follow once we prove the reverse inclusion $\mathcal{F}_{\text{supp}}(\pi_1 S) \subset \mathcal{F}_{\text{supp}}[\partial S]$, which we do using Stallings' method from [Sta00]. That inclusion is obvious if $\mathcal{F}_{\text{supp}}[\partial S] = \{[F_n]\}$, so suppose that $\mathcal{F}_{\text{supp}}[\partial S]$ is a proper free factor system of F_n . Let L be an F_n -marked graph having a proper subgraph K with noncontractible components such that $[\pi_1 K] = \mathcal{F}_{\text{supp}}[\partial S]$.

Consider the π_1 -injective map $f: S \rightarrow L$ obtained by composing $S \xrightarrow{j} Y = X \xrightarrow{d} G \mapsto L$ where the final map is a homotopy equivalence preserving marking. We may homotope this map so that $f(\partial S) \subset K$. By general position we may perturb f to be transverse to the set M consisting of the midpoints of the edges of $L \setminus K$. By π_1 -injectivity of f , each component

of $f^{-1}(M)$ is a circle bounding a disc in the interior of S . Choosing an innermost such disc D , and using that G is an Eilenberg-Maclane space, by a further homotopy supported on a neighborhood of D we may alter the map so as to remove ∂D from the inverse image of M . By induction, we may assume $f(S) \cap M = \emptyset$. By a further homotopy rel ∂S we may push S entirely into K . Since S is connected, we have pushed it into a single component of K , from which it follows that $\mathcal{F}_{\text{supp}}[\pi_1 S] \sqsubset [\pi_1 K] = \mathcal{F}_{\text{supp}}[\partial S]$.

Item (2) is a consequence of Fact 1.42 and Definition 2.4 (4a) which together imply that $[\partial_0 S]$ is represented by a height r circuit in $G = G_r$. Here is another proof that avoids Fact 1.42, again based on Stallings method. Arguing by contradiction, if $\mathcal{F}_{\text{supp}}(\partial_0 S) \sqsubset [G_{r-1}]$ then the deformation retraction $d: X \rightarrow G$ may be homotoped relative to G_{r-1} so that it takes $\partial_0 S$ to G_{r-1} . Again we may assume that d is transverse to the set M of midpoints of edges of $G \setminus G_{r-1}$. Each component of $f^{-1}(M)$ is an embedded circle in $\text{int}(S)$ which is homotopically trivial in X . By π_1 -injectivity of $j: S \rightarrow X$, each such circle is homotopically trivial in S and so bounds a disc in $\text{int}(S)$. Using this fact and proceeding just as above, we may alter d by homotopy rel G_{r-1} so that $d(S) \subset G_{r-1}$, contradicting that $d: S \rightarrow G$ is a homotopy equivalence.

To prove (3), invariance of $[\pi_1 S]$ and of $[\partial S]$ is built into Definition 2.4, and so their free factor supports are invariant as well. From (2) and Definition 2.4 it follows that $[\partial_i S]$ is supported by $[\pi_1 G_{r-1}]$ if and only if $i \neq 0$, and together with ϕ -invariance of $[\pi_1 G_{r-1}]$ it follows that ϕ preserves $[\partial_0 S]$ and the set $\{[\partial_1 S], \dots, [\partial_m S]\}$, and hence their free factor supports are preserved as well.

To prove (4), it suffices by Fact 1.24 (3) to prove that there does not exist any ϕ -invariant free factor system \mathcal{F} for which there are strict inclusions

$$\mathcal{F}_{\text{supp}}\{[\partial_1 S], \dots, [\partial_m S]\} \sqsubset \mathcal{F} \sqsubset \mathcal{F}_{\text{supp}}[\pi_1 S] \quad (*)$$

By Theorem 1.30 we may assume that the CT $f: G \rightarrow G$ has a strict inclusion of core filtration elements $G_s \subset G_t \subset G_u$ realizing the free factor systems $(*)$, respectively. We may then assume as above that $Y = Y^0 = X$ and $G_r = G_r^0 = G$. From (2) it follows that $[\pi_1 S] \not\sqsubset [\pi_1 G_{r-1}]$ and so $G_u = G_r = G$, $\mathcal{F}_{\text{supp}}[\pi_1 S] = \{[F_n]\}$, and $G_s \subset G_t \subset G_{r-1}$. For each $i = 1, \dots, m$ we have $[\partial_i S] \sqsubset \mathcal{F}_{\text{supp}}(G_s)$ and so the attaching map $\gamma_i: \partial_i S \hookrightarrow G_{r-1}$ is freely homotopic to a circuit in G_s given by an immersion $\hat{\gamma}_i: \partial_i S \rightarrow G_{r-2}$. Define a 2-complex \hat{Y} and a quotient map $\hat{j}: G_{r-1} \amalg S \rightarrow \hat{Y}$ by re-attaching $\partial_i S$ to G_{r-1} using the maps $\hat{\gamma}_i$ instead of the maps γ_i for $i = 1, \dots, m$. We obtain a homotopy equivalence $\hat{h}: \hat{Y} \rightarrow Y$ which restricts to the identity on G_{r-1} and such that the maps $S \xrightarrow{j} Y$ and $S \xrightarrow{\hat{j}} \hat{Y} \xrightarrow{\hat{h}} Y$ are homotopic. Letting $\Sigma \subset S$ be a spine which contains each $\partial_i S$, $i = 1, \dots, m$, there exist deformation retractions $S \mapsto \Sigma$ and $\hat{d}: \hat{Y} \rightarrow G_{r-1} \cup \hat{j}(\Sigma)$ such that the composition $S \mapsto \Sigma \xrightarrow{\hat{j}} G_s \cup \hat{j}(\Sigma) \subset G_{r-1} \cup \hat{j}(\Sigma)$ equals the composition $S \xrightarrow{j} \hat{Y} \xrightarrow{\hat{d}} G_{r-1} \cup \hat{j}(\Sigma)$. We may regard $G_{r-1} \cup \hat{j}(\Sigma)$ as a marked graph via the homotopy equivalence $G_{r-1} \cup \hat{j}(\Sigma) \hookrightarrow \hat{Y} \xrightarrow{\hat{h}} Y \xrightarrow{d} G$. The precomposition of this homotopy equivalence with the map $S \mapsto G_{r-1} \cup \hat{j}(\Sigma)$ is homotopic to the map $d|S: S \rightarrow G$, and it follows that we have an inclusion of subgroup systems $[\pi_1 S] \sqsubset \mathcal{F}_{\text{supp}}(G_s \cup \hat{j}(\Sigma))$. But $G_s \cup \hat{j}(\Sigma)$ is a proper subgraph of $G_{r-1} \cup \hat{j}(\Sigma)$ because it does not contain any edge of $G_t \setminus G_s$, and it follows that $\mathcal{F}_{\text{supp}}[\pi_1 S]$ is a proper free factor system, contradicting that $\mathcal{F}_{\text{supp}}[\pi_1 S] = \{[F_n]\}$. \square

2.1.4 Characterizing geometric strata: Proof of Fact 2.3

The “furthermore” clause follows from Fact 1.41.

The implication (1) \implies (2) for “improved relative train track representatives” is contained in [BFH00], Theorem 5.1.5, item eg-(iii), and the proof of this particular item is found on page 590 of [BFH00].

In our present CT context, we prove (1) \implies (2) as follows. Assuming H_r to be geometric, consider a weak geometric model Y as notated in Definition 2.2, with quotient map $j: G_{r-1} \amalg S \rightarrow Y$. Since $\partial_0 S$ is preserved by the pseudo-Anosov homeomorphism $\Psi: S \rightarrow S$, it follows that the circuit $c = d_\#(\partial_0 S)$ is invariant by $f_\#$. By Fact 1.39, the circuit c splits completely into fixed edges and indivisible Nielsen paths. By Lemma 2.5 (2) the circuit c has height r , and so one of the terms of the complete splitting of c must have height r , but that term cannot be a fixed edge, so by Fact 1.40 that term is the unique (up to reversal) indivisible Nielsen path ρ_r of height r . The path ρ_r must be closed for otherwise, by Fact 1.42 (1b), one of the endpoints of ρ_r is not contained in G_{r-1} , but then the term of the complete splitting of c incident to that endpoint is an edge of H_r and so is not a fixed edge nor an indivisible Nielsen path, a contradiction.

Remark. Observe that this proof of (1) \implies (2) does not make use of the fact that $d: Y \rightarrow G_r$ is a deformation retraction, only that it is a homotopy equivalence which restricts to the inclusion $G_{r-1} \subset G_r$. As such, this proof works just as stated using instead the definition of geometricity given in [BFH00] Definition 5.1.4, and so that definition also implies item (2). Combining this with the proofs to come of (2) \implies (3) \implies (1), we obtain a proof that geometricity as defined in [BFH00] Definition 5.1.4 implies geometricity as defined here in Definition 2.2.

The implication (2) \implies (3) is the “only if” direction of Fact 1.42 (2).

We now prove (3) \implies (1). So, assume that there exists an indivisible Nielsen path ρ_r of height r crossing each edge of H_r exactly twice. By Fact 1.41 (1) and Fact 1.44, each component of G_{r-1} is noncontractible, a fact which we assume henceforth without comment. Let G_r^0 be the component of G_r containing H_r , and hence also containing ρ_r . It suffices to construct a weak geometric model Y^0 of the restricted CT $f|G_r^0$ for its stratum H_r , because that automatically extends to a weak geometric model Y of $f|G_r$ for H_r by extending the homotopy equivalence of Y^0 to a homotopy equivalence of Y which, on $Y - Y^0 = G - G_r^0$, is equal to f . We may therefore assume that $G = G_r^0 = G_r$.

In order to construct Y we shall follow the method laid out in the proof of Proposition 5.3.1 of [BFH00]. All of the hypotheses of that proposition are satisfied except that we replace the hypothesis “ f is \mathcal{F} -Nielsen minimized” by the statement (EG Nielsen Paths) of Definition 1.29. Lemma 4.18 of [FH11] says that with this replacement, the statement and proof of Proposition 5.3.1 of [BFH00] remain valid. As said earlier, that proof mostly produces a weak geometric model, but among other things we must carefully note the steps of that proof in which item (4) of Definition 2.1 is established. For that purpose we go through some details of the proof of [BFH00] Proposition 5.3.1, carefully citing various steps.

Following [BFH00], page 571 in the proof of Proposition 5.3.1, let $Y = M(\rho_r)$ be the

mapping cylinder of $\rho_r: [0, 1] \rightarrow G_r$, which recall is the target of the quotient map

$$q: Q \amalg G_r \rightarrow M(\rho_r), \quad Q = [0, 1] \times [0, 1]$$

that identifies $(s, 0) \in Q$ to $\rho_r(s) \in G_r$ for each $s \in [0, 1]$. The embedding $G_r \rightarrow Y = M(\rho_r)$ is simply the restriction of q . Denote the “upper boundary” of Q to be $\partial_u Q = (\{0\} \times [0, 1]) \cup ([0, 1] \times \{1\}) \cup (\{1\} \times [0, 1])$. Choose a deformation retraction $d_Q: Q \rightarrow [0, 1] \times \{0\}$ that restricts to a homeomorphism $\partial_u Q \rightarrow [0, 1] \times \{0\}$. Clearly d_Q induces a deformation retraction $d: M(\rho_r) = Y \rightarrow G_r$, and on Q we have the commutative relation $d \circ q = q \circ d_Q$ (by construction, $d \circ q | \partial_u Q$ is a parameterization of the Nielsen path ρ_r). Note that already we have an embedding $G_r \hookrightarrow Y$ and deformation retraction $d: Y \rightarrow G_r$ satisfying item (4b).

Next we define the surface S and its boundary circles $\partial_0 S, \dots, \partial_m S$ required for item (1) of Definition 2.1. Subdivide $[0, 1] \approx [0, 1] \times 0 \subset Q$ into subintervals each mapped by ρ_r to an edge of G_r , so each edge of H_r occurs exactly twice. Define a quotient map $q_S: Q \rightarrow S$ by identifying in pairs, consistent with ρ_r , those subintervals of $[0, 1] \times 0$ which are mapped to edges of H_r , and so S is a compact surface with nonempty boundary. Clearly $q_S(\partial_u Q)$ is a connected subset of ∂S and so is contained in a component which we define to be the upper boundary $\partial_0 S$; the remaining components, if any, are enumerated arbitrarily as $\partial_1 S, \dots, \partial_m S$.

Next we define the quotient map $j: S \amalg G_{r-1} \rightarrow Y$ referred to in item (3) of Definition 2.1, although we do not yet specify exactly the identifications made by this quotient map that are needed in order to completely verify (3). The quotient map q clearly factors into two quotient maps as follows:

$$Q \amalg G_r \xrightarrow{q_S} S \amalg G_r \xrightarrow{j_r} M(\rho_r) = Y$$

Consider the 1-complex $\kappa = q_S([0, 1] \times 0) \subset S$. The deformation retraction $d_Q: Q \rightarrow [0, 1] \times 0$ clearly induces a deformation retraction $S \rightarrow \kappa$ whose homotopy inverse is the inclusion $\kappa \hookrightarrow S$. The edges of κ whose interiors are contained in the interior of S are labelled one-to-one by edges of H_r . The remaining edges of κ form a subcomplex κ_0 labelled (not necessarily one-to-one) by edges of G_{r-1} . Clearly we have

$$\kappa_0 = \partial S \setminus q_S(\partial_u Q)$$

Under the embedding of G_r in $M(\rho_r)$, the restriction $j_r | \kappa$ may be identified with the labelling map $\gamma: \kappa \rightarrow G_r$, the image of which contains H_r . By restricting j_r we there obtain a surjective map $j: S \amalg G_{r-1} \rightarrow Y$, and clearly this is a quotient map.

We verify further items of Definition 2.1 with citations from the proof of Proposition 5.3.1 on p. 578 of [BFH00], several of which are applications of [BFH00] Lemma 5.3.9, using the fact that the relative train track map $f | G_r$ is reduced; in our context this fact is part of property (Filtration) in the definition of a CT.

Next we verify (4a) of Definition 2.1. The proof of Proposition 5.3.1 uses Lemma 5.3.9 to show that the quotient map q_S identifies the endpoints of the arc $\partial_u Q$, and so we obtain $\partial_0 S$ as the quotient of $\partial_u Q$ under that identification. We have already seen that the quotient map $q: Q \amalg G_r \rightarrow Y$ restricts on $\partial_u Q$ to a parameterization of ρ_r , and so it follows that the quotient map $j: S \amalg G_{r-1} \rightarrow Y$ restricts on $\partial_0 S$ to a parameterization of ρ_r . It is also

clear from the construction that $G_r \cap \partial_0 S$ is the base point p_r of ρ_r , and we have verified Definition 2.1 (4a).

Next, for $i = 1, \dots, m$ we define α_i and prove item (2) of Definition 2.1. At this stage we know that $\kappa_0 = \partial S \setminus \partial_0 S$ is the union of the lower boundary components $\partial_1 S, \dots, \partial_m S$, and on each of these we define $\alpha_i = \gamma | \partial_i S: \partial_i S \rightarrow G_{r-1}$. The proof of Proposition 5.3.1 uses Lemma 5.3.9 to show that α_i is homotopically nontrivial.

Next we verify item (3) of Definition 2.1. We have already verified that the quotient map $j: S \amalg G_{r-1} \rightarrow Y$ does indeed identify each $x \in \partial_i S$ to $\alpha_i(x)$, for each $i = 1, \dots, m$. We must verify that j identifies no other point pairs. By construction, the only point pairs of $S \amalg G_{r-1}$ that are identified by j are point pairs in $G_{r-1} \cup \kappa$. Also, j maps κ_0 to G_{r-1} and maps the union of the interiors of the edges of $\kappa - \kappa_0$ injectively to the union of the interiors of the edges of H_r , which are disjoint from G_{r-1} . Letting V be the set of vertices of κ not in κ_0 , we therefore need only consider point pairs in $G_{r-1} \cup \kappa_0 \cup V$. Also, the only identifications of point pairs of $G_{r-1} \cup \kappa_0$ made by j are those made by the union of the maps $\alpha_i: \partial_i S \rightarrow G_{r-1}$, which are defined just by restricting j to the components of $\kappa_0 = \partial_1 S \cup \dots \cup \partial_m S$. It therefore only remains to consider two points of V or a point of V and a point of G_{r-1} . We again cite the proof of Proposition 5.3.1, which uses Lemma 5.3.9 once again to show that for each $v \in V$ we have $j_r(\text{Link}(\kappa, v)) = \text{Link}(G_r, j_r(v))$, which implies that $j(v) = j_r(v) \notin G_{r-1}$, and that if $v \neq w \in V$ then $j(v) = j_r(v) \neq j_r(w) = j(w)$.

We turn to item (5) of Definition 2.1. Defining the homotopy equivalence $h: Y \rightarrow Y$ to be the composition $Y \xrightarrow{d} G \xrightarrow{f} G \hookrightarrow Y$, this clearly satisfies the dynamic Definition 2.1 (5a). It remains to construct a homeomorphism $\Psi: (S, \partial_0 S) \rightarrow (S, \partial_0 S)$ with pseudo-Anosov mapping class and to verify item (5b) of Definition 2.1. These tasks are accomplished in the final portions of the proof of Proposition 5.3.1 on p. 579 of [BFH00] to which we make one last citation. In outline, the brunt of that proof is contained in Corollary 5.3.8 of [BFH00] which says that the homotopy equivalence inclusion $\kappa \hookrightarrow S$, when composed with the map $S \xrightarrow{j} Y \xrightarrow{d} G \xrightarrow{f} G$ may be lifted to a homotopy equivalence $f_\kappa: \kappa \rightarrow \kappa$ that permutes the free homotopy classes of the components of κ_0 . One also checks that f_κ fixes the free homotopy class of the circuit of κ that is the image of $\partial_0 S$ under the deformation retraction $S \rightarrow \kappa$. Conjugating f_κ by the homotopy equivalence between S and κ one obtains a homotopy equivalence of S preserving the free homotopy classes of ∂S , in particular preserving the class of $\partial_0 S$. By the Dehn-Nielsen-Baer theorem [FM12] this map is homotopic the desired homeomorphism $\Psi: (S, \partial_0 S) \rightarrow (S, \partial_0 S)$. A further check that this map has no periodic conjugacy classes other than the boundary components implies that its mapping class is pseudo-Anosov.

This completes the proof of Fact 2.3.

Remark. Unlike in the proof of Proposition 5.3.1, we do not in the end replace α_i with an immersion α'_i in the same free homotopy class, because this would destroy the mapping cylinder properties that are used to obtain the embedding $G_r \hookrightarrow Y$ and the deformation retraction $d: Y \rightarrow G_r$. In this regard, on the bottom of page 578 of [BFH00] the proof of Proposition 5.3.1 incorrectly asserts that the mapping cylinder is homeomorphic to the space obtained by gluing S to Y using the immersions α'_i as gluing maps, but this assertion is inconsequential for anything else in [BFH00].

2.2 Complementary subgraph and peripheral splitting

For this section we fix a rotationless $\phi \in \text{Out}(F_n)$, a representative CT $f: G \rightarrow G$, a geometric EG-stratum H_r , and a geometric model X for H_r with associated notation from Definition 2.4.

Definition 2.6 (The complementary subgraph). Define the *complementary subgraph* of X to be

$$L = \text{cl}(X - \text{int}(S)) = (G \setminus H_r) \cup \partial_0 S$$

This graph L is naturally identified with the quotient of the disjoint union of $G \setminus H_r$ and an arc E_ρ modulo the following identifications: if $p_r \in G \setminus H_r$ —equivalently p_r is not an interior point of G_r —then both endpoints of E_ρ are identified to p_r ; otherwise, the endpoints of E_ρ are identified just with each other, forming a circle component of L . In either case, E_ρ forms a loop in L based at p_r that is identified with $\partial_0 S$. Item (4a) of Definition 2.4 implies that the composition $E_\rho \hookrightarrow (G \setminus H_r) \amalg E_\rho \hookrightarrow L \hookrightarrow X \xrightarrow{d} G$ is a parameterization of the Nielsen path ρ . Definition 2.4 (4) implies that $L \cap \text{int}(S)$ is just the set of attaching points of X .

Listing the noncontractible components of L as $\{L_l\}$, associated to L is the subgroup system in $\pi_1(X) \approx F_n$ defined by $[\pi_1 L] = \{[d_* \pi_1(L_l)]\}$.

Lemma 2.7 addresses malnormality, or failure thereof, for the subgroup systems of F_n associated to the complementary subgraph L and the surface S . The proof will be an application of Bass-Serre theory, depending on the construction the “peripheral splitting” of F_n that is associated to X . This is a certain graph of groups presentation of F_n obtained by splitting X at the components of ∂S and the attaching points. In Section 2.6 we shall also use Bass-Serre theory to characterize the “free boundary circles” of X .

Lemma 2.7 will be used in several later arguments in this section, including the proofs of Propositions 2.11 and 2.14.

Lemma 2.7. *The maps $L \xrightarrow{d} X$ and $S \xrightarrow{j} X$ are π_1 -injective (for the latter see also the discussion at the end of Definition 2.4). Furthermore:*

- (1) *The subgroup system $[\pi_1 L]$ in $\pi_1(X) \approx F_n$ is malnormal and has one component for each noncontractible component L_l , that is, if $l \neq l'$ then $[d_* \pi_1 L_l] \neq [d_* \pi_1 L_{l'}]$. As a consequence:*
 - (a) *Distinct circuits in L map via d to distinct circuits in G .*
 - (b) *For any subgraph $K < L$ with noncontractible components $\{K_j\}$, each map $K_j \hookrightarrow G \hookrightarrow X$ is π_1 -injective, the subgroup system $[\pi_1 K] = \{[d_* \pi_1 K_j]\}$ in $\pi_1(X) \approx F_n$ is malnormal, and if $j \neq j'$ then $[d_* \pi_1 K_j] \neq [d_* \pi_1 K_{j'}]$.*
- (2) *The subgroup $d_* \pi_1 S < F_n$ is its own normalizer (although it need not be malnormal). For any nontrivial $c \in \pi_1(S) \hookrightarrow F_n$ with conjugacy class $[c]$ in F_n , if $[c]$ is carried by $[\pi_1 L]$, or if there exists $c' \in \pi_1(S)$ such that c, c' are conjugate in F_n but not in $\pi_1(S)$, then c is peripheral in $\pi_1(S)$ meaning that c is represented by a loop in ∂S .*

For an example where $\pi_1 S$ is not malnormal, let S have three boundary components, and attach the two lower boundary components to the same circuit in G_{r-1} .

To define the peripheral splitting of F_n we shall use the method of [SW79] by first constructing a graph of spaces, whose underlying 2-complex \widehat{X} is obtained from X by blowing up the attaching points in a certain manner, as follows.

Definition 2.8 (Blowing up attaching points). Let the attaching points of X be enumerated as $\{z_k\}_{k=1}^q$. Let $v_k \geq 1$ be the valence of z_k in the graph L . Choose a regular neighborhood U of the finite set $\{z_k\}$, with component U_k containing z_k . We have $U_k = (U_k \cap \text{int}(S)) \cup (U_k \cap L)$ where the subset $U_k \cap \text{int}(S)$ is an open disc in $\text{int}(S)$ containing z_k , the subset $U_k \cap L$ is a union of v_k half-open arcs with common endpoint z_k , and the intersection of those two subsets is the point z_k . Let \widehat{X} be the 2-complex obtained from X by pulling each attaching point z_k apart, replacing U_k with the disjoint union of the two sets $U_k \cap \text{int}(S)$ and $U_k \cap L$ together with an arc η_k having one endpoint z_k^S identified with the copy of z_k in $U_k \cap \text{int}(S)$ and having opposite endpoint z_k^L identified with the copy of z_k in $U_k \cap L$. We note that the quotient map $\widehat{X} \rightarrow X$ obtained by collapsing each η_k to the point z_k is a homotopy equivalence, and so there is an induced isomorphism $\pi_1(\widehat{X}) \rightarrow \pi_1(X) \approx F_n$ (well defined, as usual, up to inner automorphism).

Recall that a *graph of spaces* structure on \widehat{X} is defined by exhibiting \widehat{X} as a certain quotient of a collection of path connected *edge spaces* and *vertex spaces*. Each edge space is exhibited as the product of a compact arc and space called the *core* of the edge space; we shall refer to the endpoints of the arc crossed with the core as the two *endpoint cores* of that edge space, and the interior points of the arc crossed with the core as the *interior cores*. The entire graph of spaces is obtained from the disjoint union of the vertex spaces and the edge spaces by attaching each boundary core to some vertex space via some π_1 -injective map. Each edge space minus its two boundary cores forms an *open edge space* which is foliated by interior cores and which maps homeomorphically via the quotient to an open subset of the graph of spaces.

Definition 2.9 (The peripheral graph of spaces). Define the “peripheral graph of spaces” structure on \widehat{X} is defined as follows. The components of ∂S are enumerated as $\partial_i S$, $i = 0, \dots, m$. Let N be a closed regular neighborhood of ∂S in the surface S , chosen so that N is disjoint from the closure of U . The component of N containing $\partial_i S$ is an annulus N_i exhibited as the product of a compact arc and a circle. Let $\widehat{S} = \text{int}(S) - \cup_{i=0}^m \text{int}(N_i)$, a compact surface with boundary which is a deformation retract of $\text{int}(S)$ and of S itself. The vertex spaces of \widehat{X} are \widehat{S} and the components of L . Each annulus N_i , $i = 0, \dots, m$ is an edge space with core being a circle: one endpoint core of N_i is attached to \widehat{S} by the inclusion map $\partial N_i \cap \widehat{S} \hookrightarrow \widehat{S}$, which is π_1 -injective because \widehat{S} is not a disc; the other endpoint core of N_i is attached to some component of L by the restricted map $\partial N_i \cap \partial S \xrightarrow{j} L$, and π_1 -injectivity follows from Definition 2.1 (??). Also, each closed arc η_k is an edge space with core being a point, and with endpoint cores attached to $z_k^S \in \widehat{S}$ and $z_k^L \in L$.

Definition 2.10 (The peripheral splitting of F_n associated to X). This is the graph of groups presentation $\Gamma(X)$ of $F_n \approx \pi_1(\widehat{X})$ that is associated to the graph of spaces structure on \widehat{X} , as explained in [SW79]. In detail, the underlying graph of $\Gamma(X)$ is the quotient of

\widehat{X} obtained by collapsing each vertex space to a vertex of $\Gamma(X)$ and collapsing each edge space to an edge of Γ by collapsing each of its interior cores to a point. The graph $\Gamma(X)$ has one “ S -vertex” obtained by collapsing \widehat{S} , and a finite set of “ L -vertices” obtained by collapsing the components of L . Also $\Gamma(X)$ has one edge for each N_i and one for each z_k . The graph $\Gamma(X)$ is bipartite: each edge has one endpoint on the \widehat{S} vertex and one endpoint amongst the L vertices. By choosing a base point in each vertex space and in the core of each edge space, and then taking fundamental groups, we associate a vertex group to each vertex of $\Gamma(X)$ and an edge group to each edge; in particular, the group of the S -vertex is $\pi_1(S)$, the group of each L -vertex is π_1 of the associated component of L , the group of each N_i edge space is infinite cyclic, and the group of each z_k edge space is trivial. By appropriate choice of paths connecting up base points, we associate to each endpoint of each edge a monomorphism from the edge group of that edge to the vertex group of that endpoint. The fundamental group $\pi_1(\Gamma(X))$ is defined as in [Ser80] or in [SW79], as is the isomorphism $\pi_1(\Gamma(X)) \approx \pi_1(\widehat{X}) \approx F_n$, well defined up to inner automorphism.

Let $T(X)$ be the Bass-Serre tree of $\Gamma(X)$, on which $\pi_1(\Gamma(X)) \approx F_n$ acts with quotient graph of groups $\Gamma(X) \approx T(X)/F_n$. A vertex of $T(X)$ is referred to as an L -vertex or S -vertex according to which its image under the orbit map $T(X) \mapsto \Gamma(X)$ is an L -vertex or S -vertex.

Proof of Lemma 2.7. π_1 -injectivity of the maps $L \rightarrow G \hookrightarrow X$ and $S \hookrightarrow X$ follows directly from the general fact that vertex groups of any graph of groups inject into its fundamental group.

To prove (1), consider two L -vertices $V \neq W$ of $T(X)$. We must prove that $\text{Stab}(V) \cap \text{Stab}(W)$ is trivial. Since $\Gamma(X)$ and hence $T(X)$ is bipartite, the arc \overline{VW} contains some S -vertex U . Letting $E \neq E'$ be the edges of \overline{VW} incident to U it follows that $\text{Stab}(V) \cap \text{Stab}(W) < \text{Stab}(E) \cap \text{Stab}(E')$, and so it suffices to prove that $\text{Stab}(E) \cap \text{Stab}(E')$ is trivial. Suppose it is nontrivial. The surface \widehat{S} has a hyperbolic structure which arises as follows: for some properly discontinuous, cocompact action $\pi_1(\widehat{S}) \curvearrowright \mathbf{H}^2$ whose limit set is a $\pi_1(\widehat{S})$ -invariant Cantor set $C \subset \partial \mathbf{H}^2$, the group $\pi_1(\widehat{S})$ acts on the convex hull $\mathcal{H}(C) \subset \mathbf{H}^2$, and \widehat{S} is identified isometrically with the quotient $\mathcal{H}(C)/\pi_1(\widehat{S})$. There exist geodesic lines $L \neq L'$ which are components of $\partial \mathcal{H}(C)$ such that under the action of $\pi_1(\widehat{S})$ on $\mathcal{H}(C)$ we have $\text{Stab}(E) = \text{Stab}(L)$ and $\text{Stab}(E') = \text{Stab}(L')$. But $\text{Stab}(L) \cap \text{Stab}(L')$ is clearly trivial, because the lines L, L' have disjoint endpoint pairs in C . Item (1a) is an immediate consequence. Item (1b) follows by using the fact that for each component L_l of L , the components of K in L_l define a malnormal subgroup system in $\pi_1(L_l)$.

To prove (2), fix an S -vertex V of $T(X)$. The identification of $T(X)/F_n$ with $\Gamma(X)$ induces an isomorphism between $\text{Stab}(V)$ and the image of the injection $\pi_1(S) \approx \pi_1(\widehat{S}) \hookrightarrow \pi_1(X) \approx F_n$. Consider another vertex $W \neq V$ of $T(X)$ such that $\text{Stab}(V) \cap \text{Stab}(W)$ is nontrivial. Letting E be the first edge of the segment from V to W we have $\text{Stab}(V) \cap \text{Stab}(W) \subset \text{Stab}(E)$. As shown above, either $\text{Stab}(E)$ is trivial or $\text{Stab}(E) = \text{Stab}(L)$ for some boundary line L of the universal cover of \widehat{S} , and so any element of $\text{Stab}(V) \cap \text{Stab}(W)$ is peripheral in $\pi_1(S)$. If W is also an S -vertex then $\text{Stab}(V) \neq \text{Stab}(W)$ which implies that $\pi_1(S)$ is its own normalizer in F_n . Letting $c \in \pi_1(S) \approx \text{Stab}(V)$ be as in (2), the desired conclusion follows once the correct W is chosen: if $[c]$ is carried by $[\pi_1 L]$ then take W to be the unique L -vertex such that $c \in \text{Stab}(W)$; and if c is conjugate in F_n to $c' \in \pi_1(S)$ but c, c' are not conjugate in $\pi_1(S)$ then take W to be the S -vertex $W = \theta \cdot V$ where $c' = \theta c \theta^{-1}$. \square

2.3 The laminations of a geometric stratum

The main result of this section is Proposition 2.11 which says that the lamination pair of a geometric stratum can be identified with the stable/unstable geodesic lamination pair of the associated pseudo-Anosov mapping class. This result can be bound in [BFH00], embedded in the proof of the geometric case of Proposition 6.0.8. We give a different proof, based on little bit of Nielsen–Thurston theory for surface mapping classes, and we get a little bit more out of the proof by identifying the free factor supports of these laminations with the free factor support of the surface S in the geometric model.

2.3.1 Review of Nielsen–Thurston theory.

We begin with a review of basic facts about geodesic laminations and pseudo-Anosov mapping classes; in the proof of Proposition 2.11 we will quote various other facts about the Nielsen–Thurston theory. The main references for this theory are [CB88] and [Mil82]. Our concern here is *solely* with non-closed surfaces, and we will state results only in that case. We shall take the liberty to state definitions and results in a manner that illuminates the connection with geometric strata and laminations in free groups; the reader may easily compare these statements with the standard forms found in the references.

In this section let \mathbf{H}^2 denote the hyperbolic plane and S_∞^1 its boundary at infinity, identified with the Gromov boundary.

Gromov compactification. Let S be a compact surface with nonempty boundary ∂S . We assume that $\chi(S) \leq -1$, and so $\pi_1(S)$ is a free group of rank ≥ 2 with Gromov boundary denoted $\partial\pi_1(S)$. Any geodesic metric on S lifts to a geodesic metric on the universal cover \tilde{S} , the identity map on \tilde{S} is a quasi-isometry between any two such lifted metrics, and for any such metric and any $x_0 \in \tilde{S}$ the orbit map $\pi_1(S) \rightarrow \tilde{S}$ defined by $\gamma \mapsto \gamma \cdot x_0$ is also a quasi-isometry. It follows that there is a well-defined *Gromov compactification* $\tilde{S} \cup \partial\pi_1(S)$ such that the actions of $\pi_1(S)$ on \tilde{S} by deck transformations and on $\partial\pi_1(S)$ piece together to give an action by homeomorphisms on $\tilde{S} \cup \partial\pi_1(S)$.

Hyperbolic structures. By a *hyperbolic structure* on S we shall mean a hyperbolic metric on S with totally geodesic boundary. Associated to a hyperbolic structure μ is its *holonomy* which is a properly discontinuous action $\rho_\mu: \pi_1(S) \rightarrow \text{Isom}(\mathbf{H}^2)$ well defined up to conjugacy by an element of $\text{Isom}(\mathbf{H}^2)$. The holonomy ρ_μ has a limit set $\text{Lim}_\mu \subset S_\infty^1 = \partial\mathbf{H}^2$ which is well-defined as the accumulation set of any orbit. The convex hull of the limit set $\mathcal{H}(\text{Lim}_\mu)$ is the smallest closed convex subset of \mathbf{H}^2 containing each geodesic $(\xi, \eta) \subset \mathbf{H}^2$ such that $\xi \neq \eta \in \text{Lim}_\mu$; we also denote $[\xi, \eta] = (\xi, \eta) \cup \{\xi, \eta\} \subset \mathbf{H}^2 \cup S_\infty^1$. Also associated to μ is its *developing map*, $D^\mu: \tilde{S} \rightarrow \mathcal{H}(\text{Lim}_\mu)$ which is a $\pi_1(S)$ equivariant isometry from \tilde{S} , with hyperbolic metric lifted from S , to $\mathcal{H}(\text{Lim}_\mu)$, with hyperbolic metric restricted from \mathbf{H}^2 , so that D^μ conjugates the deck transformation action of $\pi_1 S$ on \tilde{S} to the action ρ_μ on $\mathcal{H}(\text{Lim}_\mu)$. There exists a unique equivariant continuous homeomorphism $D_\infty^\mu: \partial\pi_1 S \rightarrow \text{Lim}_\mu$, and this homeomorphism pieces together with D^μ to produce an equivariant homeomorphism defined on the Gromov compactification

$$D^\mu \amalg D_\infty^\mu: \tilde{S} \amalg \partial\pi_1 S \rightarrow \mathcal{H}(\text{Lim}_\mu) \amalg \text{Lim}_\mu$$

When μ has been specified we often use this homeomorphism to identify the domain and range.

For any $\tilde{\ell} \in \widetilde{\mathcal{B}}(\pi_1 S)$ its *realization* $\tilde{\ell}_\mu \subset \mathcal{H}(\text{Lim}_\mu) \subset \mathbf{H}^2$ is the hyperbolic geodesic connecting the two points of $\tilde{\ell}$. For any $\ell \in \mathcal{B}(\pi_1 S)$, its realization ℓ_μ is the bi-infinite geodesic in S obtained by choosing any lift $\tilde{\ell} \in \widetilde{\mathcal{B}}(\pi_1 S)$ and projection $\tilde{\ell}_\mu$ to S ; this is well-defined independent of the choice.

Geodesic laminations. A *geodesic lamination* on S is a closed subset $\Lambda \subset \mathcal{B}(\pi_1(S))$ such that for some (any) hyperbolic structure μ on S the following hold: for each $\ell \in \Lambda$ the corresponding geodesic ℓ_μ on S is either a bi-infinite simple geodesic or it wraps infinite-to-one around a simple closed geodesic on S ; if $\ell \neq \ell' \in \Lambda$ then $\ell_\mu \cap \ell'_\mu = \emptyset$; and $\Lambda_\mu = \cup_{\ell \in \Lambda} \ell_\mu$ is a compact subset of $\text{int}(S)$ which we refer to as the *realization* of Λ in the hyperbolic structure μ . In the literature, the set $\Lambda_\mu \subset S$ itself is referred to as a geodesic lamination on S with respect to μ , but for present purposes we take the liberty of focussing on $\Lambda \subset \mathcal{B}(\pi_1 S)$.

A geodesic lamination Λ is *minimal* if every leaf of Λ is dense in Λ , from which it follows that every leaf is birecurrent. If Λ is minimal then either Λ is a single periodic leaf or no leaf of Λ is periodic. Assuming Λ is minimal, we say that Λ *fills* S if for some (any) finite geodesic metric μ , and for each component R of $S - \Lambda$, one of the following holds: R is disjoint from ∂S and R is contractible; or R contains a unique component of ∂S and R deformation restricts to that component. Note that the filling property automatically precludes the possibility that Λ is a single periodic leaf.

Pseudo-Anosov mapping classes. The *mapping class group* of S is

$$\mathcal{MCG}(S) = \text{Homeo}(S)/\text{Homeo}_0(S)$$

where $\text{Homeo}_0(S)$ is the normal subgroup of the homeomorphism group $\text{Homeo}(S)$ consisting of those homeomorphisms that are isotopic to the identity. By the Dehn-Nielsen-Baer theorem [FM12], the kernel of the natural map $\text{Homeo}(S) \rightarrow \text{Out}(\pi_1(S))$ is $\text{Homeo}_0(S)$ and its image is the subgroup $\text{Out}(\pi_1 S, \partial S) < \text{Out}(\pi_1(S))$ that preserves the set of conjugacy classes of cyclic subgroups associated to the components of ∂S , and so we obtain a natural isomorphism $\mathcal{MCG}(S) \rightarrow \text{Out}(\pi_1 S, \partial S)$.

Given a mapping class $\phi \in \mathcal{MCG}(S)$ and a minimal lamination $\Lambda \subset \mathcal{B}(\pi_1(S))$ which is not a single periodic leaf, and given a number $\lambda > 1$, suppose that $\phi(\Lambda) = \Lambda$ and suppose that there exists a $\pi_1(S)$ -equivariant, positive Borel measure μ on the lifted lamination $\widetilde{\Lambda} \subset \widetilde{\mathcal{B}}(\pi_1(S))$ and an automorphism $\Phi \in \text{Aut}(\pi_1(S))$ representing ϕ such that $\Phi_*(\mu) = \lambda\mu$. If $\lambda > 1$ then we say that Λ is an *unstable lamination* for ϕ with expansion factor λ ; whereas if $\lambda < 1$ then Λ is a *stable lamination* for ϕ with contraction factor λ . A leaf ℓ of Λ^s or Λ^u is *fixed* if $\phi(\ell) = \ell$ and is *periodic* if $\phi^k(\ell) = \ell$ for some $k \neq 0$.

A mapping class $\phi \in S$ is *pseudo-Anosov* if there exists a pair of minimal, filling geodesic laminations $\Lambda^s \neq \Lambda^u \subset \mathcal{B}(\pi_1(S))$ such that Λ^s is a stable lamination and Λ^u an unstable lamination for ϕ . In this situation ϕ determines the laminations Λ^s , Λ^u uniquely and they are called the *stable* and *unstable* laminations of ϕ . Also, ϕ uniquely determines the expansion factor for Λ^u and the contraction factor for Λ^s , and those numbers are multiplicative inverses $\lambda > 1$ and $\lambda^{-1} < 1$, respectively.

2.3.2 Comparing free group laminations and surface laminations.

Consider a rotationless $\phi \in \text{Out}(F_n)$ represented by a CT $f: G \rightarrow G$ with geometric EG stratum H_r . Consider also a geometric model X with weak geometric submodel Y of $f: G \rightarrow G$ relative to H_r , with all associated notations from Definitions 2.1 and 2.4, in particular the surface S and pseudo-Anosov mapping class $\psi \in \mathcal{MCG}(S)$.

By π_1 -injectivity of the map $S \xrightarrow{j} Y \hookrightarrow X$ combined with the fact that finite rank subgroups of F_n are quasiconvex, we obtain a $\pi_1 S$ -equivariant embedding $\partial(\pi_1 S) \hookrightarrow \partial F_n$. As explained in Section 1.2.1, that embedding induces a $\pi_1 S$ -equivariant embedding $\tilde{\mathcal{B}}(\pi_1 S) \subset \tilde{\mathcal{B}}$ which induces in turn a continuous function $\mathcal{B}(\pi_1 S) \rightarrow \mathcal{B}$.

Proposition 2.11. *With the notation above, letting $\Lambda^+ \in \mathcal{L}(\phi)$ be the lamination associated to H_r and $\Lambda^- \in \mathcal{L}(\phi^{-1})$ its dual lamination, and letting $\Lambda^s, \Lambda^u \subset \mathcal{B}(S)$ be the stable and unstable laminations of the pseudo-Anosov mapping class $\psi \in \mathcal{MCG}(S)$ associated to the geometric model, we have:*

- (1) *The map $\mathcal{B}(\pi_1 S) \rightarrow \mathcal{B}$ takes Λ^u, Λ^s homeomorphically to Λ^+, Λ^- respectively.*
- (2) *Every leaf of Λ^+ is dense in Λ^+ , and similarly for Λ^- .*
- (3) *All leaves of Λ^+ and Λ^- are generic.*
- (4) $\mathcal{F}_{\text{supp}}(\Lambda^+) = \mathcal{F}_{\text{supp}}(\Lambda^-) = \mathcal{F}_{\text{supp}}[\pi_1 S]$.

Proof. Note immediately that (1) \implies (2), and that (2) \implies (3) using that each leaf of a minimal lamination is birecurrent. It remains to prove (1) and (4).

We first prove (1) and (4) in the special case that $X = Y = S$ and so $F_n = \pi_1(S)$. We need the fact that periodic leaves of Λ^u exist, which is an easy consequence of the Nielsen–Thurston theory. One proof uses that in any hyperbolic structure the realization Λ_μ^u has only finitely many boundary leaves, and this set of leaves is invariant under the action of ϕ (see [CB88]). Another proof uses that any pseudo-Anosov homeomorphism representing ψ has a dense set of periodic points, and so its unstable measured foliation has periodic leaves (see [FLP⁺79]), and so Λ^u has periodic leaves.

Consider now a periodic leaf ℓ of Λ^u and choose a lift $\tilde{\ell} \subset \tilde{\Lambda}^u$. Choose $k \geq 1$ so that $\phi^k(\ell) = \ell$. We may then choose $\Phi' \in \text{Aut}(\pi_1(S))$ to be an automorphism representing ϕ^k and fixing ℓ . Let $\ell = \{\xi, \eta\} \subset \tilde{\mathcal{B}}(\pi_1(S))$. After possibly replacing k by $2k$ we may assume $\Phi'(\xi) = \xi, \Phi'(\eta) = \eta$. Fixing a finite hyperbolic structure μ and identifying $\partial\pi_1(S)$ with Lim_μ , by applying Nielsen Theory there exist disjoint neighborhoods $U_\xi, U_\eta \subset \partial\pi_1(S)$ of ξ, η which are attracting neighborhoods for ξ and η , e.g. for ξ this means that $U_\xi \supset \Phi'(U_\xi) \supset (\Phi')^2(U_\xi) \supset \dots$ and this sequence is a neighborhood basis for ξ in $\partial\pi_1(S)$. From this it follows that ℓ has an attracting neighborhood under the action of ϕ^k on $\mathcal{B}(\pi_1(S))$, namely, the image in $\mathcal{B}(\pi_1(S))$ of all lines in $\tilde{\mathcal{B}}(\pi_1(S))$ having one endpoint in U_ξ and the other endpoint in U_η . Since Λ^u is minimal, ℓ is birecurrent. Since Λ^u is filling, the realization of ℓ in any finite hyperbolic structure does not cover a closed geodesic. It follows that ℓ satisfies the definition of a generic leaf of an attracting lamination of ψ . Since Λ^u is minimal, it follows that Λ^u is the weak closure of ℓ , and so $\Lambda^u \in \mathcal{L}(\psi)$. The same argument applied to the inverse shows that $\Lambda^s \in \mathcal{L}(\psi^{-1})$.

Next, working in the finite rank free group $\pi_1(S)$, we show that $\mathcal{F}_{\text{supp}}(\Lambda^u) = \{[\pi_1(S)]\}$; applying this to ψ^{-1} it follows that $\mathcal{F}_{\text{supp}}(\Lambda^s) = \{[\pi_1(S)]\}$. The free factor system $\mathcal{F}_{\text{supp}}(\Lambda^u)$ is connected, and so by Lemma 2.5 (1) it suffices to prove that for any connected free factor system $\{[A]\}$ of $\pi_1(S)$ that supports Λ^u , and for any component b of ∂S , $[A]$ supports b . Pick a finite hyperbolic structure μ and let R_b be the component of $S - \Lambda_\mu$ containing b , so R_b is an open subset of S . As shown for example in [CB88], the completion \overline{R}_b of the hyperbolic metric on R_b is a crown surface, homeomorphic to $S^1 \times [0, 1] - A \times 1$ for some finite subset $A \subset S^1$, such that each component of $(S^1 - A) \times 1$ is a leaf of Λ_μ ; let L be the number of such leaves. Letting \tilde{R}_b be a connected lift of R_b in the universal cover $\text{Lim}_\mu \approx \tilde{S}$, the stabilizer Z_b of \tilde{R}_b is an infinite cyclic subgroup of the deck transformation group. Also, $\partial \tilde{R}_b$ consists of a bi-infinite geodesic \tilde{b} covering b and stabilized by Z_b , and a bi-infinite sequence of oriented leaves ℓ_i of $\tilde{\Lambda}_\mu$, $i \in \mathbb{Z}$, such that the terminal endpoint of ℓ_i equals the initial endpoint of ℓ_{i+1} in Lim_μ . Furthermore, we can choose a generator z for Z_b such that $z(\ell_i) = \ell_{i+L}$. Choose a free factor A representing the conjugacy class $[A]$ so that $\partial \tilde{\ell}_0 \subset \partial A$. By combining malnormality of A with Fact 1.2 and an induction argument, it follows that $\partial \tilde{\ell}_i \subset \partial A$ for all $i \in \mathbb{Z}$. The collection of points $\cup_{i \in \mathbb{Z}} \partial \tilde{\ell}_i$ is invariant under Z_b and so they accumulate on the two points of $\partial \tilde{b}$, and it follows that $\partial \tilde{b} \subset \partial A$, implying that b is carried by $[A]$ as desired.

We now turn to the general case of (1) and (4). We use the notation established in setting up Proposition 2.11. By appropriate choices of base points and paths amongst them, the composition $S \xrightarrow{j} Y \hookrightarrow X$ induces an injection

$$\pi_1(S) \hookrightarrow \pi_1(Y) \approx \pi_1(G_r) \subset \pi_1(G) \approx F_n \quad (*)$$

We identify $\pi_1(S)$ with its image in F_n under the injection $(*)$, and we let $\beta: \mathcal{B}(\pi_1 S) \rightarrow \mathcal{B}(F_n)$ be the induced continuous map. From the dynamic conditions in Definitions 2.4 and 2.1 it follows that we may choose a homeomorphism $\Psi: S \rightarrow S$ representing the mapping class $\psi \in \mathcal{MCG}(S)$, and we may choose an automorphism $\Phi \in \text{Aut}(F_n)$ representing ϕ which leaves $\pi_1 S$ invariant, so that the automorphism $\Phi \mid \pi_1 S$ is a representative of the pseudo-Anosov mapping class ψ under the injection $\mathcal{MCG}(S) \approx \text{Out}(\pi_1 S, \partial S) \hookrightarrow \text{Out}(\pi_1 S)$. Applying the special case combined with Lemma 1.61, it follows that $\beta(\Lambda^u)$, $\beta(\Lambda^s)$ are a dual lamination pair in $\mathcal{L}^\pm(\phi)$ whose free factor support in F_n equals $\mathcal{F}_{\text{supp}}[\pi_1 S]$. By Lemma 2.5 (2) we have $\mathcal{F}_{\text{supp}}[\pi_1 S] \not\subset [G_{r-1}]$, but from $(*)$ we have $\mathcal{F}_{\text{supp}}[\pi_1 S] \subset [G_r]$. It follows that $\mathcal{F}_{\text{supp}}(\beta(\Lambda^u)) \not\subset [G_{r-1}]$ and $\mathcal{F}_{\text{supp}}(\beta(\Lambda^u)) \subset [G_r]$. But the only element of $\mathcal{L}(\phi)$ with this property is Λ^+ , and so $\beta(\Lambda^u) = \Lambda^+$. By duality it follows that $\beta(\Lambda^s) = \Lambda^-$. This proves (4) and also proves most of (1).

What is left is to prove that the maps $\Lambda^u \rightarrow \Lambda^+$ and $\Lambda^s \rightarrow \Lambda^-$ obtained by restriction $\beta: \mathcal{B}[\pi_1 S] \rightarrow \mathcal{B}$ are homeomorphisms. We prove this for Λ^u , the same following for Λ^s by replacing ϕ with ϕ^{-1} . The idea of the proof is that although $\pi_1 S < F_n$ need not be a free factor, nor even malnormal, by Bass-Serre theory it just misses being malnormal, in that the intersections of $\pi_1 S$ with any of its conjugates by elements of $F_n - \pi_1 S$ are peripheral in $\pi_1 S$. Since no leaf of Λ^u is peripheral, the map β restricts to a bijection between Λ^u and Λ^+ . In the Hausdorff setting this would be enough to verify that this bijection is a homeomorphism, but we must work a little harder in the non-Hausdorff setting.

We may assume henceforth that $G = G_r$ and that $X = Y$ and so we may identify

$\pi_1(G_r) \approx \pi_1(Y) \approx F_n$; this assumption is justified by restricting $f: G \rightarrow G$ to the component G_r^0 of G_r that contains H_r , and by restricting the geometric model Y for H_r to the component Y_r^0 of Y that contains H_r .

It follows by Lemma 2.7 (2) that if $\gamma \in F_n - \pi_1(S)$ then $\pi_1(S) \cap \pi_1(S)^\gamma$ is either the trivial subgroup or a peripheral subgroup in $\pi_1(S)$ meaning an infinite cyclic subgroup conjugate into the fundamental group of some component of ∂S . Applying this together with Fact 1.2 it follows that for any $\gamma \in F_n - \pi_1(S)$ the intersection $\partial\pi_1(S) \cap \gamma \cdot \partial\pi_1(S)$ is either empty or is the pair of points fixed by the stabilizer of some component of $\partial\widetilde{S}$. This implies in turn that $\widetilde{\mathcal{B}}(\pi_1 S) \cap \gamma \cdot \widetilde{\mathcal{B}}(\pi_1 S)$ is either empty or a single line corresponding to a component of $\partial\widetilde{S}$. Let Γ be a set of left coset representatives of $\pi_1(S)$ in F_n , and let $\gamma_0 \in \pi_1(S) \cap \Gamma$ be the representative of the coset $\pi_1(S)$. Applying Fact 1.8 it follows that the subset $\cup\{\gamma \cdot \widetilde{\mathcal{B}}(\pi_1 S) \mid \gamma \in \Gamma - \gamma_0\}$ is closed in $\widetilde{\mathcal{B}}$. Also, the finite subset $\mathcal{B}^\partial(\pi_1 S) \subset \mathcal{B}(\pi_1 S)$ represented by components of ∂S is closed in $\mathcal{B}(\pi_1 S)$, and so its total lift $\widetilde{\mathcal{B}}^\partial(\pi_1 S)$ consisting of all lines corresponding to components of $\partial\widetilde{S}$ is closed in $\widetilde{\mathcal{B}}(\pi_1 S)$ and so is also closed in $\widetilde{\mathcal{B}}$. The set

$$\mathcal{B} - \left(\left(\bigcup_{\gamma \in \Gamma - \gamma_0} \gamma \cdot \widetilde{\mathcal{B}}(\pi_1 S) \right) \cup \widetilde{\mathcal{B}}^\partial(\pi_1 S) \right)$$

is therefore open in \mathcal{B} . The intersection of this set with $\widetilde{\mathcal{B}}[\pi_1 S] = \cup\{\gamma \cdot \widetilde{\mathcal{B}}(\pi_1 S) \mid \gamma \in \Gamma\}$ is equal to $\widetilde{\mathcal{B}}^{\text{int}}(\pi_1 S) = \widetilde{\mathcal{B}}(\pi_1 S) - \widetilde{\mathcal{B}}^\partial(\pi_1 S)$ and so the latter is open in $\widetilde{\mathcal{B}}[\pi_1 S]$. Furthermore, $\widetilde{\mathcal{B}}^{\text{int}}(\pi_1 S)$ is disjoint from each of its translates $\gamma \cdot \widetilde{\mathcal{B}}^{\text{int}}(\pi_1 S)$ for $\gamma \in \Gamma - \gamma_0$. The collection of sets $\{\gamma \cdot \widetilde{\mathcal{B}}^{\text{int}}(\pi_1 S) \mid \gamma \in \Gamma\}$ is therefore pairwise disjoint, and each set in this collection is an open subset of the union of the collection.

Consider now the lamination $\Lambda^u \subset \mathcal{B}(\pi_1 S)$ and its image Λ^+ under the continuous map $\beta: \mathcal{B}(\pi_1 S) \rightarrow \mathcal{B}$. By definition of unstable laminations we have $\Lambda^u \subset \mathcal{B}^{\text{int}}(\pi_1 S)$. Let $\widetilde{\Lambda}^u \subset \widetilde{\mathcal{B}}(\pi_1 S)$ be its total lift with respect to $\pi_1 S$ and so $\widetilde{\Lambda}^u \subset \widetilde{\mathcal{B}}^{\text{int}}(\pi_1 S)$. The total lift of $\widetilde{\Lambda}^u$ with respect to $\pi_1 S$ is equal to the set $\cup_{\gamma \in \Gamma} (\gamma \cdot \widetilde{\Lambda}^u)$, and from what is proved in the previous paragraph this is a disjoint union whose terms are each open in the union. It follows that the inclusion of $\widetilde{\Lambda}^u$ into the union $\cup_{\gamma \in \Gamma} (\gamma \cdot \widetilde{\Lambda}^u)$ descends to a homeomorphism from $\Lambda^u = \widetilde{\Lambda}^u / \pi_1 S$ to $\Lambda^+ = (\cup_{\gamma \in \Gamma} \gamma \cdot \widetilde{\Lambda}^u) / F_n$, but this homeomorphism is precisely the restriction to Λ^u of the map $\beta: \mathcal{B}(\pi_1 S) \rightarrow \mathcal{B}$.

This completes the proof of Proposition 2.11. □

2.3.3 Application: The inclusion lattice of attracting laminations.

The inclusion partial order on subsets of \mathcal{B} restricts to an inclusion partial order on $\mathcal{L}(\phi)$. An attracting lamination is said to be *topmost* if it is maximal with respect to inclusion on $\mathcal{L}(\phi)$. It was proved in [BFH00] Corollary 6.0.1 that a lamination in $\mathcal{L}(\phi)$ is topmost if and only if its dual lamination in $\mathcal{L}(\phi^{-1})$ is topmost. We extend this result to show that the duality relation respects inclusion; the proof uses Proposition 2.11.

Lemma 2.12. *If Λ_i^\pm and Λ_j^\pm are dual lamination pairs for $\phi \in \text{Out}(F_n)$ then $\Lambda_i^+ \subset \Lambda_j^+$ if and only $\Lambda_i^- \subset \Lambda_j^-$.*

Proof. Passing to a power we may assume ϕ is rotationless. Let F^i, F^j denote free factors with associated one component free factor systems:

$$\begin{aligned}\{[F^i]\} &= \mathcal{F}_{\text{supp}}(\Lambda_i^+) = \mathcal{F}_{\text{supp}}(\Lambda_i^-) \\ \{[F^j]\} &= \mathcal{F}_{\text{supp}}(\Lambda_j^+) = \mathcal{F}_{\text{supp}}(\Lambda_j^-)\end{aligned}$$

If $[F^j]$ does not carry $[F^i]$ then $\Lambda_j^+ \not\supset \Lambda_i^+$ and $\Lambda_j^- \not\supset \Lambda_i^-$. We may therefore assume that $[F^j]$ carries $[F^i]$. Restricting ϕ to F^j we may assume that $F^j = F_n$, which implies that Λ_j^+ is topmost (in $\mathcal{L}(\phi)$) and Λ_j^- is topmost (in $\mathcal{L}(\phi^{-1})$). By Corollary 6.0.11 of [BFH00] Λ_i^+ is topmost if and only if Λ_i^- is topmost. In this topmost case, $\Lambda_j^+ \not\supset \Lambda_i^+$ and $\Lambda_j^- \not\supset \Lambda_i^-$. We may therefore assume that neither Λ_i^+ nor Λ_i^- is topmost, in which case it suffices to show that $\Lambda_i^+ \subset \Lambda_j^+$ and that $\Lambda_i^- \subset \Lambda_j^-$; the two cases are similar, so it suffices to assume that $\Lambda_j^+ \not\supset \Lambda_i^+$ and argue to a contradiction.

Since Λ_i^+ is not topmost and is not contained in Λ_j^+ , there is a topmost lamination $\Lambda_k^+ \in \mathcal{L}(\phi)$ such that (A) $\Lambda_i^+ \subset \Lambda_k^+$ and (B) $\Lambda_k^+ \not\subset \Lambda_j^+$. Choose a CT $f: G \rightarrow G$ representing ϕ and let H_r and H_s be the EG strata corresponding to Λ_i^+ and Λ_k^+ respectively. By (A), Λ_k^+ has nongeneric lines. By Proposition 2.11 (3), the stratum of G corresponding to the lamination Λ_k^+ is non-geometric. By (B), Λ_j^+ is not weakly attracted to Λ_k^+ , that is, no generic leaf of Λ_j^+ is weakly attracted to a generic leaf of Λ_k^+ . Applying the Weak Attraction Theorem 6.0.1 and Remark 6.0.2 of [BFH00] to the topmost lamination Λ_k^+ it follows that Λ_j^+ is carried by a proper free factor of F_n , which contradicts our assumption that $F^j = F_n$. \square

2.4 Geometricity is an invariant of a dual lamination pair

Given $\phi \in \text{Out}(F_n)$, the following proposition shows that geometricity is an invariant of a lamination $\Lambda \in \mathcal{L}(\phi)$, not just an invariant of a particular CT representing ϕ and EG stratum corresponding to Λ . This is accomplished by exhibiting a property of Λ that is evidently independent of the choice of a CT, namely condition (2). Furthermore this property shows that geometricity is an invariant of dual lamination pairs $\Lambda^\pm \in \mathcal{L}^\pm(\phi)$, not just of the two individual laminations. Because of this we may extend the terminology “geometric” beyond the setting of EG strata to apply to those elements of $\mathcal{L}(\phi)$ and of $\mathcal{L}^\pm(\phi)$ that satisfy property (2).

Proposition 2.13. *For any $\phi \in \text{Out}(F_n)$, any $\Lambda_\phi^\pm \in \mathcal{L}^\pm(\phi)$, and any CT $f: G \rightarrow G$ representing a rotationless iterate of ϕ , if $H_r \subset G$ is the EG stratum associated to Λ_ϕ^+ then the following are equivalent:*

- (1) H_r is a geometric stratum.
- (2) There exists a finite ϕ -invariant set $\mathcal{C} = \{[c_1], \dots, [c_k]\}$ of conjugacy classes such that $\mathcal{F}_{\text{supp}}(\mathcal{C}) = \mathcal{F}_{\text{supp}}(\Lambda_\phi^\pm)$.

Furthermore, if (1), (2) hold then for any CT $f': G' \rightarrow G'$ representing a positive or negative rotationless power of ϕ the following holds:

- (3) If $H'_s \subset G'$ is the EG stratum associated to Λ_ϕ^+ or Λ_ϕ^- then H'_s is geometric.
- (4) If $[G_r] = [G'_s]$ and $[G_{r-1}] = [G'_{s-1}]$, and if ρ_r, ρ'_s are the closed indivisible Nielsen paths in G, G' of height r, s , respectively, then the conjugacy classes $[\rho_r], [\rho'_s]$ are the same up to inverse.
- (5) Adopting the notation of the geometric model for H_r , we have

$$\mathcal{F}_{\text{supp}}[\partial S] = \mathcal{F}_{\text{supp}}[\pi_1 S] = \mathcal{F}_{\text{supp}}(\Lambda_\phi^\pm)$$

Proof. Item (3) follows from the equivalence of (1) and (2), the fact that $\mathcal{F}_{\text{supp}}(\Lambda_\phi^+) = \mathcal{F}_{\text{supp}}(\Lambda_\phi^-)$, and the fact that (2) is independent of the choice of positive or negative power of ϕ and choice of CT $f: G \rightarrow G$ representing that power. To prove (4), by Fact 1.42 the conjugacy classes $[\rho_r], [\rho'_s]$ have the same characterization and are therefore equal (up to inverse): it is the unique (up to inverse) root-free conjugacy class that is carried by $[G_r] = [G'_s]$ but not by $[G_{r-1}] = [G'_{s-1}]$ and is fixed by ϕ .

To prove (2) \implies (1), suppose H_r is not geometric. By Fact 1.42, no conjugacy class of height r is fixed by ϕ , and so every fixed conjugacy class of height $\leq r$ is supported by $[G_{r-1}]$; however, Λ_ϕ^+ itself is not supported by $[G_{r-1}]$.

Finally, note that (1) \implies (5) follows from Lemma 2.5 and Proposition 2.11 (4), and (5) \implies (2) using $\mathcal{C} = [\partial S]$. \square

2.5 Stabilizing its complement also stabilizes the surface

We continue with the fixed notation of Section 2.2: a rotationless $\phi \in \text{Out}(F_n)$, a representative CT $f: G \rightarrow G$, a geometric stratum H_r , and a geometric model X for H_r , and all the notations of Definition 2.4 associated to X . We also adopt the notation of the complementary subgraph $L \subset X$ from Definition 2.6, and the notation of the peripheral graph of spaces \widehat{X} from Definition 2.9.

Consider a closed oriented surface S , and a curve system $C \subset S$ separating S into two components whose closures are V, W , neither of which is a disc. Any homeomorphism of S that stabilizes V up to isotopy must also stabilize W up to isotopy, and vice versa. It follows that subgroups $\text{Stab}[V], \text{Stab}[W] < \mathcal{MCG}(S) = \text{Out}(\pi_1 S)$ that stabilize the isotopy classes of V, W , respectively, are equal. This kind of thing fails in $\text{Out}(F_n)$ when one tries to use a free factorization $F_n = A * B$ in place of the decomposition $S = V \cup W$: an outer automorphism that preserves the conjugacy class $[A]$ need not preserve the conjugacy class B , and vice versa, so neither of the subgroups $\text{Stab}[A], \text{Stab}[B]$ is contained in the other.

There is however an $\text{Out}(F_n)$ context involving geometric models in which we do obtain complementary behavior of this sort. Recalling from Section 1.1.2 the stabilizer of a subgroup system, the following result says that the subgroup $\text{Stab}[\pi_1 L] < \text{Out}(F_n)$ is contained in the subgroup $\text{Stab}[\pi_1 S] < \text{Out}(F_n)$, in fact it is contained in the $\mathcal{MCG}(S)$ subgroup of $\text{Stab}[\pi_1 S]$, and that we obtain a well-defined induced homomorphism $\text{Stab}[\pi_1 L] \hookrightarrow \mathcal{MCG}(S)$. We will need this statement in greater generality, replacing L by certain subgraphs $K \subset L$.

For the statement of the theorem, recall that by the Dehn-Nielsen-Baer Theorem [FM12], the subgroup of $\text{Out}(\pi_1 S)$ consisting of all $\theta \in \text{Out}(\pi_1 S)$ that preserve the set of conjugacy classes associated to all oriented components of ∂S is naturally isomorphic to $\mathcal{MCG}(S)$, and we identify $\mathcal{MCG}(S)$ with this subgroup.

Proposition 2.14. *For any subcomplex $K \subset L$ such that $j(\partial S) \subset K$, and for any $\theta \in \text{Out}(F_n)$ such that $\theta[\pi_1 K] = [\pi_1 K]$, we have:*

- (1) θ fixes $[\pi_1 S]$.
- (2) *The restriction $\theta \mid \pi_1 S \in \text{Out}(\pi_1 S)$ is well-defined, it lives in $\mathcal{MCG}(S)$, and the induced function $\text{Stab}[\pi_1 K] \mapsto \mathcal{MCG}(S)$ defined by $\theta \mapsto \theta \mid \pi_1 S$ is a homomorphism.*

The proof of Proposition 2.14, most of which is in the following lemma, takes up the rest of the section. Once item (1) is proved, the part of item (2) saying that $\theta \mid \pi_1 S \in \text{Out}(\pi_1 S)$ is well-defined follows from Lemma 2.7 (2) and Fact 1.4; then, once it is proved that $\theta \mid \pi_1 S$ lives in $\mathcal{MCG}(S)$, the rest of item (2) follows from Fact 1.4. Item (1) itself, and the fact that $\theta \mid \pi_1 S$ lives in $\mathcal{MCG}(S)$, are each an immediate consequence of the following:

Lemma 2.15. *Under the same hypothesis as Proposition 2.14, and letting K' be the union of the noncontractible components of K , for any $\theta \in \text{Out}(F_n)$ such that $\theta[\pi_1 K] = [\pi_1 K]$ there exists a homotopy equivalence of pairs $\Theta_X: (X, K') \rightarrow (X, K')$ and a homeomorphism $\Theta_S: S \rightarrow S$ such that:*

- (1) Θ_X represents θ with respect to the isomorphism $\pi_1 X \xrightarrow{d_*} \pi_1(G) \approx F_n$.
- (2) *The maps $j \circ \Theta_S$ and $\Theta_X \circ j: S \rightarrow X$ are homotopic.*

Proof. Since the components of ∂S are π_1 -injective in X , the j -image of each is contained in a component of K' .

List the components of K' as $K_1 \cup \dots \cup K_\ell$. From Lemma 2.7, the subgroup system $[\pi_1 K]$ has exactly ℓ distinct components $[\pi_1 K_1], \dots, [\pi_1 K_\ell]$, and if $i \neq j \in \{1, \dots, \ell\}$ then no nontrivial conjugacy class of F_n is carried by both $[\pi_1 K_i]$ and $[\pi_1 K_j]$.

First we construct $\Theta_X: (X, K') \rightarrow (X, K')$ satisfying (1). Start with any homotopy equivalence $\Theta: X \rightarrow X$ that represents θ . For each component K_i of K' there exists another component $K_{\sigma(i)}$ such that $\Theta_*[\pi_1 K_i] = [\pi_1 K_{\sigma(i)}]$. The function $\sigma: \{1, \dots, \ell\} \rightarrow \{1, \dots, \ell\}$ must be a permutation, because otherwise circuits in distinct conjugacy classes will have images in the same conjugacy class. Since X and each K_i are Eilenberg-Maclane spaces, we can homotope each of the restrictions $\Theta \mid K_i$ to a map $\Theta_i: K_i \rightarrow K_{\sigma(i)}$ which is a π_1 -isomorphism, and each such Θ_i is a homotopy equivalence. The union of functions $\Theta_1 \cup \dots \cup \Theta_\ell: K' \rightarrow K'$ is therefore also a homotopy equivalence. By the homotopy extension property, the homotopy from $\Theta \mid K'$ to $\Theta_1 \cup \dots \cup \Theta_\ell$ extends to a homotopy of all of Θ .

Next we construct the homeomorphism $\Theta_S: S \rightarrow S$ satisfying (2). For this purpose we arrange to improve the map $g_1 = \Theta_X \circ j: S \rightarrow X$ by a sequence of homotopies. To start, we already have $g_1(\partial S) = \Theta_X(j(\partial S)) \subset \Theta(K) \subset K$.

By homotoping g_1 rel ∂S , we shall produce a map $g_2: S \rightarrow X$ such that $g_2(S) \subset j(S) \cup K$, still retaining the property that $g_2(\partial S) \subset K$. We again use Stallings method from [Sta00] as in the proof of Lemma 2.5. Consider $X \setminus (j(S) \cup K)$ which is a union $E_1 \cup \dots \cup E_\nu$ of “naked” edges of X , meaning edges of L whose interiors are open subsets of X . Letting M denote the set of midpoints of the edges E_1, \dots, E_ν , we may perturb g_1 to be transverse to M , preserving that $g_1(\partial S) \subset K$. The set $g_1^{-1}(M)$ is disjoint union of circles each of which, by π_1 -injectivity, bounds a disc in the interior of S . Starting with an innermost disc

D , by a homotopy of g_1 supported on a slightly larger disc we may remove ∂D from $g_1^{-1}(M)$, and after finitely many of these homotopies we have arranged that $g_1(S) \cap M = \emptyset$. We may then post-compose g_1 with a deformation retraction $(X \setminus (j(S) \cup K)) - M \mapsto j(S) \cup K$, to produce the desired map $g_2: S \rightarrow X$.

By the Simplicial Approximation Theorem [Spa81] we may subdivide S and X as simplicial complexes, and perturb g_2 by a homotopy, preserving the property that $g_2(\partial S) \subset K$, so that g_2 is a simplicial map. Having done that, consider the subcomplex $R \subset S$ defined by $R = g_2^{-1}(L) = g_2^{-1}(K)$, where this equation follows from the fact that $g_2(S) \subset j(S) \cup K$ and the fact that $j(S) \cap L \subset K$. Note in particular that $\partial S \subset R$.

We claim that any (homotopically) nontrivial closed curve γ in S that is contained in R is peripheral in S , meaning that γ is homotopic to a closed curve in ∂S . To see why, note that $\Theta(j(\gamma))$ is a nontrivial closed curve in X homotopic to the closed curve $g_2(\gamma)$ in K . Since Θ restricts to a homotopy equivalence of K , there exists a nontrivial closed curve γ' in K such that $\Theta(\gamma')$ is homotopic to $\Theta(j(\gamma))$. Since Θ is a homotopy equivalence of X , the closed curves γ' and $j(\gamma)$ are homotopic. Applying Lemma 2.7 (2) it follows that γ is peripheral in S , proving the claim.

Next, by further homotopy of g_2 , we produce a map $g_3: S \rightarrow X$ such that $g_3(S) \subset j(\text{int}(S)) = X - L$. Let $N(R)$ be a regular neighborhood of R in the surface S , and so $N(R)$ is a compact subsurface of S containing ∂S . We shall prove that $N(R)$ can be engulfed in a regular neighborhood of ∂S . From the claim of the previous paragraph it follows that any closed curve contained in $N(R)$, in particular any component of $\partial N(R)$, is either homotopically trivial or peripheral in S . Let $N^+(R)$ be the union of $N(R)$ with every disc bounded by a homotopically trivial component of $\partial N(R)$ and every annulus bounded by two peripheral components of $\partial N(R)$. Again $N^+(R)$ is a compact subsurface of S containing R and so containing ∂S , but now each component of $N^+(R)$ is either a disc or an annulus neighborhood of ∂S . Let D_1, \dots, D_J be the disc components of $N^+(R)$. Let $\alpha_1, \dots, \alpha_J$ be a pairwise disjoint collection of arcs properly embedded in the subsurface $S - N^+(R)$ such that α_j has one endpoint on ∂D_j and another endpoint on an annulus component of $N^+(R)$. Let $N^{++}(R)$ be a regular neighborhood of $N^+(R) \cup (\alpha_1 \cup \dots \cup \alpha_J)$. It follows that $N^{++}(R)$ is a disjunction union of annulus neighborhoods of the components of ∂S , that is, a regular neighborhood of ∂S . By the Regular Neighborhood Theorem, $N^{++}(R)$ is ambiently isotopic in S to any regular neighborhood, and so the surface S is therefore isotopic in itself to its subsurface $\text{cl}(S - N^{++}(R))$. Postcomposing this isotopy with the map g_2 , the resulting composed map $g_3: S \rightarrow X$ given by $S \mapsto \text{cl}(S - N^{++}(R)) \subset S \xrightarrow{g_2} X$ has image contained in $X - L = j(\text{int}(S))$, as desired.

Since $g_3(S) \subset j(\text{int}(S))$ and since j restricts to an embedding on $\text{int}(S)$, we may lift g_3 to a self map $\tilde{g}_3: S \rightarrow S$. By construction it follows that the maps $j \circ \tilde{g}_3$ and $\Theta_X \circ j: S \rightarrow X$ are homotopic, and so it suffices to prove that \tilde{g}_3 is homotopic to a homeomorphism of S .

The induced homomorphism of fundamental groups $(\tilde{g}_3)_*: \pi_1(S) \rightarrow \pi_1(S)$ is an automorphism, because with respect to the embedding $\pi_1(S) \rightarrow \pi_1(X)$ induced by j the map $(\tilde{g}_3)_*$ is the restriction of an automorphism of F_n representing $\theta \in \text{Out}(F_n)$, and by a theorem of Peter Scott found in Lemma 6.0.6 of [BFH00] the restriction of an automorphism of F_n to any invariant finite rank subgroup is an automorphism of that subgroup. The automorphism $(\tilde{g}_3)_*$ preserves the conjugacy classes of the boundary components of S , because by applying Lemma 2.7 (2), for any conjugacy class C in $\pi_1 S$ the following are equivalent:

C is represented by a closed curve in ∂S ; the image conjugacy class $[j(C)] = [g_3(C)]$ in F_n is carried by both $[\pi_1 K]$ and $[\pi_1 S]$. By the Dehn-Nielsen-Baer theorem [FM12], $\tilde{g}_3: S \rightarrow S$ is homotopic to a homeomorphism Θ_S . \square

2.6 Preserving the free boundary circles.

We continue with the fixed notation of the previous section.

For application in Part II [HM13b] we need Lemma 2.18 which gives a stronger conclusion than Lemma 2.15 regarding invariance of *free boundary circles* of S —by definition, these are the components $\partial_i S$ of ∂S for which there exists a half-open annulus neighborhood $A \subset S$ such that the map $S \xrightarrow{j} X$ restricts to a homeomorphism from A onto an open subset of X .

For the statement of Lemma 2.18 we will also use a stronger hypothesis than Lemma 2.15, requiring that the free factor support of the complementary subgroup system $[\pi_1 L]$ is all of F_n . From that hypothesis several consequences follow:

Fact 2.16. *If $\mathcal{F}_{\text{supp}}[\pi_1 L] = \{[F_n]\}$ then X has no attaching points and every component of L is noncontractible. Furthermore, every edge group of $\Gamma(X)$ is infinite cyclic as is every edge stabilizer of $T(X)$.*

Proof. There is a subgraph $\widehat{L} \subset \widehat{X}$ which is mapped homeomorphically to L under the homotopy equivalence quotient map $\widehat{X} \mapsto X$ that collapses each of the edges η_k to the corresponding attaching point v_k . The graph \widehat{L} is disjoint from the interior of each η_k . So if any v_k and its corresponding η_k do exist, then \widehat{L} is contained in a subcomplex of \widehat{X} which represents a proper free factor of X , contradicting that $\mathcal{F}_{\text{supp}}[\pi_1 L] = \{[F_n]\}$ and therefore proving that X has no attaching points.

To prove that L has no contractible component it suffices to prove that each valence 1 vertex $v \in L$ is contained in a noncontractible subgraph of L . First, $v \in H_r$ because if $v \notin H_r$ then, since G is a core graph, v would have valence ≥ 2 in $G \setminus H_r \subset L$. Also, $v \notin \partial_0 S$ because then v would have valence ≥ 2 in $\partial_0 S \subset L$; in particular $v \neq p_r$. Also, $v \notin \text{int}(H_r; G_{r-1})$, the interior of H_r relative to G_{r-1} , because, since $v \neq p_r$, it would follow that v is an attaching point. The only remaining possibility is that $v \in H_r - \text{int}(H_r; G_{r-1}) \subset G_{r-1}$. But each component of G_{r-1} is noncontractible (Fact 2.3).

By Definition 2.10, all edge groups are infinite cyclic or trivial, and trivial edge groups are associated only to the edges obtained by blowing up attaching points, of which there are none as we have just shown. \square

The following lemma characterizes free boundary circles in terms of the action of F_n on the Bass-Serre $T(X)$. The final characterization (4) is purely algebraic and will play an important role in proving invariance of free boundary circles.

Lemma 2.17. *For each boundary circle $\partial_i S$, letting $L_{l(i)}$ be the component of L containing $j(\partial_i S)$, letting $v \in \Gamma(X)$ be the vertex corresponding to $L_{l(i)}$, and letting $V \in T(X)$ be any vertex projecting to v , the following are equivalent:*

- (1) $\partial_i S$ is a free boundary circle.

(2) v has valence 1 in $\Gamma(X)$, and letting $e \subset \Gamma(X)$ be the edge incident to v , the injection from the edge group of e to the vertex group of v is an isomorphism.

(3) V has valence 1 in $T(X)$.

If further $\mathcal{F}_{\text{supp}}[\pi_1 L] = \{[F_n]\}$ then the following is also equivalent:

(4) For each subgroup A in the conjugacy class $[\pi_1 L_{l(i)}]$ there exists a unique subgroup B in the conjugacy class $[\pi_1 S]$ such that some finite index subgroup of A is contained in B .

Proof. The equivalence (1) \iff (2) is an immediate consequence of the construction of $\Gamma(X)$, and the equivalence (2) \iff (3) is a consequence of Bass-Serre theory.

For proving the equivalence (3) \iff (4), we may rechoose V within its orbit so that $A = \text{Stab}(V)$. The implication (3) \implies (4) holds without the “further” hypothesis, as follows. Let $V \in T(X)$ have valence 1. Letting E_1 be the unique edge of $T(X)$ incident to V and letting W_1 be the opposite vertex of E_1 , clearly W_1 projects to the \widehat{S} vertex. Setting $B = \text{Stab}(W_1)$, we have $A = \text{Stab}(V) = \text{Stab}(E_1) < \text{Stab}(W_1) = B$, and these subgroups are all nontrivial by Fact 2.16. Arguing by contradiction, if (4) does not hold then there exists a vertex $W_2 \neq W_1 \in T(X)$ also projecting to the S vertex of $\Gamma(X)$ such that the subgroup $\text{Stab}(\overline{VW_2}) = \text{Stab}(V) \cap \text{Stab}(W_2)$ is nontrivial. The arc $\overline{VW_2}$ contains the edge E_1 incident to V and at least one other edge; let E_2 be the next edge after E_1 . The subgroup $\text{Stab}(\overline{VW_2})$ is contained in the stabilizers of each of E_1, E_2 , and so $\text{Stab}(E_1) \cap \text{Stab}(E_2)$ is nontrivial. But E_1, E_2 are distinct edges incident to W_1 , and so $\text{Stab}(E_1)$ and $\text{Stab}(E_2)$ are stabilizers of distinct components of the boundary of the universal cover of \widehat{S} , which implies that $\text{Stab}(E_1) \cap \text{Stab}(E_2)$ is trivial, a contradiction.

Assume now that (4) holds. Applying (4), there is a unique vertex $W_1 \in T(X)$ projecting to the S -vertex such that the subgroup $A_1 = A \cap \text{Stab}(W_1)$ has finite index in A . Letting E_1 be the first edge on the arc $\overline{VW_1}$, it follows that $A' = A \cap B = \text{Stab}(V) \cap \text{Stab}(W_1) < \text{Stab}(E_1)$. Since L_j is noncontractible (Fact 2.16), the subgroups A and A' are nontrivial. By Fact 2.16, all edges in $T(X)$ have infinite cyclic stabilizer. Since $\text{Stab}(E_1)$ is infinite cyclic it follows that A' is infinite cyclic, and hence so is A . Assuming that V does not have valence 1, the tree $T(X)$ has another edge E_2 incident to V ; let W_2 be the vertex of E_2 opposite V . Since $\text{Stab}(E_2) \subset \text{Stab}(V) = A$, since $\text{Stab}(E_2)$ is nontrivial, and since A is infinite cyclic, it follows that $\text{Stab}(E_2)$ is infinite cyclic and has finite index in A . But since $\text{Stab}(E_2) \subset \text{Stab}(W_2)$ we have contradicted uniqueness of W_1 , proving that V has valence 1. \square

Let $\partial^f S$ be the union of free boundary circles of S . For example, $\partial_0 S \subset \partial^f S$ if and only if $p_r \notin G \setminus G_r$, equivalently p_r is an interior point of H_r ; if H_r is the top stratum then this holds automatically. The map j embeds $\partial^f S$ into X ; we identify $\partial^f S$ with its image in X . Note that $\partial^f S$ is a union of components of L ; the remaining components are denoted $L^{nf} = L - \partial^f S$. Let $\partial^{nf} S = \partial S - \partial^f S$, and notice that $j(\partial^{nf} S) \subset L^{nf}$. Recall the notation N_i for closed regular neighborhoods of the boundary circles $\partial_i S$, and let A_i be the interior of N_i which is identified by j with an open subset of X . Define the *free subsurface* S^f of S to be the closure of $S - \cup\{N_i \mid \partial_i S \subset \partial^{nf} S\}$, a compact subsurface to which S deformation retracts. The map $j: S \rightarrow X$ embeds S^f into X and we identify S^f with its image. It follows

that L^{nf} and S^f are disjoint closed sucomplexes of X and that $X - (L^{\text{nf}} \cup S^f)$ is the disjoint union of open annuli $A_i = \text{int}(N_i)$ over all i such that $\partial_i S \subset \partial^{\text{nf}} S$. Under the deformation retraction $S \mapsto S^f$ let $\partial_i S \leftrightarrow \partial_i S^f$ be the induced bijection of boundary circles: if $\partial_i S \in \partial^f S$ then $\partial_i S = \partial_i S^f$; whereas if $\partial_i S \in \partial^{\text{nf}} S$ then $\partial_i S \cap \partial_i S^f = \emptyset$ and $\partial_i S \cup \partial_i S^f = \partial N_i$. The map $j \mid N_i$ may be regarded as a homotopy in X between the embedded circle $j(\partial_i S^f) \subset X$ and the closed edge path $j(\partial_i S)$ in L^{nf} .

Lemma 2.18 (Addendum to Lemma 2.15). *For any $\theta \in \text{Out}(F_n)$ such that $\theta[\pi_1 L] = [\pi_1 L]$, and assuming that $\mathcal{F}_{\text{supp}}[\pi_1 L] = \{[F_n]\}$, there exists a homotopy equivalence $\Theta_X: (X, L) \rightarrow (X, L)$ representing θ and satisfying the conclusions of Lemma 2.15 so that the following also hold:*

- (1) *In addition to preserving L , also Θ_X preserves the sets L^{nf} , $\partial^f S$, S^f , and $X \setminus S^f$.*
- (2) *The restriction $\Theta_X \mid S^f$ is a homeomorphism of S^f .*

Proof. By Fact 2.16 the set of attaching points is empty and it follows that $L \cap j(S) = j(\partial S)$. Consider now the map $\Theta_X: (X, L) \rightarrow (X, L)$ that is produced by Lemma 2.15.

We first prove that $\Theta_X(\partial^f S) = \partial^f S$ which, by Lemma 2.17, is a union of components of L . Since each component of L is noncontractible (Fact 2.16), and since the subgroup system $[\pi_1 L]$ is malnormal with one component for each component of L (Lemma 2.7 (1)), it follows that the restriction $\Theta_X: L \rightarrow L$ permutes the components of L , that θ permutes the components of the subgroup system $[\pi_1 L]$, and that these permutations are compatible under the bijection between components of L and of $[\pi_1 L]$. Let $\alpha \in \text{Aut}(F_n)$ be any representative of θ . Since θ preserves the subgroup systems $[\pi_1 L]$ and $[\pi_1 S]$, and since their disjoint union $[\pi_1 L] \cup [\pi_1 S]$ is precisely the set of conjugacy classes of nontrivial vertex stabilizers for the action of F_n on $T(X)$, it follows that α induces a unique permutation of the vertices of $T(X)$ with nontrivial stabilizers so that $\text{Stab}(\alpha \cdot V) = \alpha(\text{Stab}(V))$, and furthermore this action preserves the L -vertices and the set of S -vertices. Given any component $\partial_i S$ of $\partial^f S$, we must find a component $\partial_j S$ of $\partial^f S$ such that $\Theta_X(\partial_i S) = \partial_j S$. Let $w \in \Gamma(X)$ be the vertex corresponding to $\partial_i S$, and let $W \in T(X)$ be any vertex projecting to w . Using that (1) \Rightarrow (4) in Lemma 2.17, there exists a unique S -vertex $V \in T(X)$ such that $\text{Stab}(W) \cap \text{Stab}(V)$ has finite index in $\text{Stab}(W)$. Since α is an automorphism of F_n , it follows that $\alpha \cdot V$ is the unique S -vertex such that $\text{Stab}(\alpha \cdot W) \cap \text{Stab}(\alpha \cdot V)$ has finite index in $\text{Stab}(\alpha \cdot W)$. Using that (4) \Rightarrow (1) in Lemma 2.17 it follows that $\alpha \cdot W$ projects to an L -vertex of $\Gamma(X)$ corresponding to some component $\partial_j S$ of $\partial^f S$, and so $\Theta_X(\partial_i S) = \partial_j S$.

We know from Lemma 2.15 that $\Theta_X \mid S^f$ is homotopic to a self-homeomorphism, and by the homotopy extension lemma that we may extend this to a homotopy of the entire map Θ_X relative to the disjoint subcomplex L^{nf} . Note that if $\partial_i S = \partial_i S^f$, $\partial_j S = \partial_j S^f$ are free boundary circles, and if the equation $\Theta_X(\partial_i S^f) = \partial_j S^f$ holds before the homotopy, then it holds as well after the homotopy. To see why, after the isotopy the right hand side of the equation is some component $\partial_k S^f$ of ∂S^f (not necessarily a free boundary circle), and it would follow that $\partial_j S^f$ and $\partial_k S^f$ are homotopic in X . But then by applying Lemma 2.7 (1) which says that $[\pi_1 L]$ is a malnormal subgroup system, and using the fact that any nonfree boundary component of S^f is homotopic into L , it follows that $\partial_j S^f = \partial_k S^f$.

Our remaining task is to homotope the map Θ_X relative to $S^f \cup L^{\text{nf}}$ so as the guarantee that Θ_X preserves $X - S^f$. Note that the components of $X - (S^f \cup L^{\text{nf}})$ are identified under

the map $j \mid \text{int}(S)$ with the *nonfree open annuli* $A_i = \text{int}(N_i)$, where the *nonfree closed annuli* $N_i \subset S$ are characterized by $\partial N_i = \partial_i S^f \cup \partial_i S$ and $\partial_i S \subset \partial^{\text{nf}} S$. It therefore suffices to consider each nonfree closed annulus N_i and the annulus map $\beta = \Theta_X \circ j: N_i \rightarrow X$, and to homotope β relative to ∂N_i so that $\beta(A_i) \subset X - S^f$.

We claim that there does exist a map $\gamma: N_i \rightarrow X$ such that $\gamma \mid \partial N_i = \beta \mid \partial N_i$, and $\gamma(A_i) \subset X - S^f$. Since N_i is just an annulus, existence of γ is equivalent to existence of a free homotopy in X between the two circle maps $\beta \mid \partial_i S^f$ and $\beta \mid \partial_i S$, such that the interior of this free homotopy is disjoint from S^f . To do this it suffices to construct appropriate free homotopies between each of these two circle maps and the circle map ζ defined to be the composition

$$\zeta: \partial_i S^f \xrightarrow{\Theta_X} \partial_k S^f \mapsto \partial_k S \xrightarrow{j} X$$

where the map $\partial_k S^f \mapsto \partial_k S$ is defined by choosing a product structure $N_k \approx S^1 \times [0, 1]$. Clearly the annulus map $N_k \xrightarrow{j} X$ itself provides a free homotopy between the circle map $\beta \mid \partial_i S^f$ and the map ζ , and the image of the interior of this free homotopy is the open annulus A_k which is disjoint from S^f . Also, the two circle maps ζ and $\beta \mid \partial_i S$ each have image in L , and they are freely homotopic in X , so by Lemma 2.7 (1a) these two circle maps are freely homotopic in L , disjoint from S^f . This proves the claim.

It remains to prove that for any two annulus maps $\beta, \gamma: N_i \rightarrow X$ which agree on ∂N_i , the map β is homotopic rel ∂N_i to a map with the same image as γ (but β need not be homotopic rel ∂N_i to γ itself). This is proved by doubling the annulus N_i across its boundary to form a torus T and defining a map $\beta \cup \gamma: T \rightarrow X$ using β on one copy of N_i and γ on the other. The image of the induced homomorphism $\mathbb{Z}^2 = \pi_1(T) \xrightarrow{(\beta \cup \gamma)_*} \pi_1(X) \approx F_n$ is infinite cyclic. Since X is an Eilenberg-Maclane space, it follows that the map $T \mapsto X$ extends continuously to a map $M \mapsto X$ defined on some solid torus M with boundary T . The two copies of N_i comprising $T = \partial M$ are homotopic rel boundary through M , and composing this homotopy with the map $M \mapsto X$ we obtain the desired homotopy rel boundary from β to a map whose image equals that of γ . \square

3 Vertex groups and vertex group systems

A free factor system in F_n consists of the conjugacy classes of nontrivial vertex stabilizers for some minimal action of F_n on a simplicial tree T such that all edge stabilizers are trivial. We generalize this by introducing the concept of a *vertex group system*, which consists of the conjugacy classes of nontrivial point stabilizers for a minimal action of F_n on some **R**-tree T with trivial arc stabilizers.²

There are two main results in this section. Proposition 3.2 is a chain condition, the proof of which was suggested to us by Mark Feighn, which bounds the length of each linear chain of vertex groups in F_n ordered by inclusion. Proposition 3.3 says that certain subgroup systems arising naturally from consideration of geometric models are vertex group systems.

²In the earlier version [HM09] “vertex group systems” were defined more broadly in terms of very small F_n -trees.

3.1 Vertex group systems

First we recall some elements of the theory of F_n -actions on \mathbf{R} -trees. An F_n tree is a minimal, isometric action $F_n \curvearrowright T$ of the group F_n on an \mathbf{R} -tree T for which no point or end of T is fixed by the whole action; we often just write T for an F_n -tree, when the action is understood. The F_n -tree is *small* is the subgroup of F_n that stabilizes any nondegenerate arc in T is either trivial or cyclic, in which case the stabilizer of each point has rank $\leq n$ [GL95, GJLL98].

A proper, nontrivial subgroup $A < F_n$ is a *vertex group* if there exists an F_n tree T with trivial arc stabilizers and an $x \in T$ such that $A = \text{Stab}(x)$, in which case we say that the vertex group A is *realized at x in T* . Each free factor is a vertex group, and is realized at some vertex of some simplicial F_n -tree with trivial edge stabilizers.

Given an F_n -tree T with trivial arc stabilizers, the *vertex group system of T* , denoted \mathcal{A}_T , is the subgroup system consisting of the conjugacy classes of all vertex groups that are realized in T . By definition of vertex group, each $A \in \mathcal{A}_T$ is proper and nontrivial. We note that in addition to the bound n on the ranks of the vertex stabilizers in T , the cardinality of \mathcal{A}_T is bounded above by a finite constant depending only on the rank n ; see [GL95, GJLL98]. The vertex group system \mathcal{A}_T is therefore an example of a subgroup system in F_n .

A subgroup system \mathcal{A} is said to be a *vertex group system* if there exists an F_n -tree T with trivial arc stabilizers such that $\mathcal{A} = \mathcal{A}_T$; we say in this situation that \mathcal{A} is *realized in T* . Note that since arc stabilizers are trivial, the point stabilized by a subgroup A in the system \mathcal{A}_T is uniquely determined by A . Each free factor system is a vertex group system, and is realized in some simplicial F_n -tree with trivial edge stabilizers.

Lemma 3.1. *A vertex group system \mathcal{A} is uniquely determined by the set \mathcal{C} of conjugacy classes of nontrivial elements of F_n that are carried by \mathcal{A} . More precisely, for any subgroup $B < F_n$ the following are equivalent:*

- (1) *The F_n -conjugacy class of each nontrivial element of B is an element of the set \mathcal{C} .*
- (2) *There exists a subgroup $A < F_n$ such that $B < A$ and $[A] \in \mathcal{A}$.*

It follows that the subgroups $A < F_n$ for which $[A] \in \mathcal{A}$ are precisely the maximal subgroups satisfying item (1).

Proof. Choose an F_n -tree T with trivial arc stabilizers in which \mathcal{A} is realized. The lemma is an immediate consequence of two elementary results in the theory of trees: \mathcal{C} consists precisely of the conjugacy classes of nontrivial elements of F_n that are elliptic on T ; and for each subgroup $B < F_n$ whose elements are all elliptic on T , there exists $p \in T$ such that $B < \text{Stab}(p)$. \square

The chain condition on vertex groups. Grushko's Theorem implies that the collection of free factors satisfies a chain condition: every free factor $A \subset F_n$ satisfies $\text{rank}(A) \leq n$, and for any free factors $A < A'$ we have $\text{rank}(A) \leq \text{rank}(A')$ with equality if and only if $A = A'$. We need the analogous result about vertex groups, the proof of which was suggested to us by Mark Feighn.

Proposition 3.2. *Every vertex group $A < F_n$ satisfies $\text{rank}(A) \leq n$. For any vertex groups $A < A'$ we have $\text{rank}(A) \leq \text{rank}(A')$ with equality if and only if $A = A'$. It follows that there is an upper bound to the length L of any properly nested sequence of vertex group systems $\mathcal{A}_1 \sqsubset \cdots \sqsubset \mathcal{A}_L$.*

Proof. The last sentence follows from the previous ones using the fact from [GL95] that gives an upper bound to the cardinality

The fact that $\text{rank}(A) \leq n$ follows from the index inequality which is the main result of [GL95]. It remains to prove that for any proper inclusion of vertex groups $A < A'$ we have $\text{rank}(A) < \text{rank}(A')$. Choose an F_n -tree T in which A is realized at some point $p \in T$. Let $S \subset T$ be a minimal A' -subtree of T , and so S is either a point or a nondegenerate subtree of T . If $p \notin S$ then A stabilizes the shortest arc of T from p to a point of S , contradicting that arc stabilizers in T are trivial. If $\{p\} = S$ then $A' < \text{Stab}(S) = \text{Stab}(p) = A$ and so $A = A'$, contradicting properness. If $p \in S$ and $\{p\} \neq S$ then S is a small A' -tree and so $\text{rank}(A) \leq \text{rank}(A')$ by the index inequality of [GL95]. If $\text{rank}(A) = \text{rank}(A')$ then by Theorem 5.2 of [Pau97], which is attributed to Levitt, it follows that S is a simplicial tree and the quotient graph of groups S/A' is a rose each edge of which is labelled by an infinite cyclic group, and so each edge of S has nontrivial stabilizer, contradicting that arc stabilizers in T are trivial. \square

Remark. In the earlier version of this paper [HM09] the reader will find Proposition 3.2 stated in the broader context of point stabilizers of very small F_n -trees. The proof, again suggested to us by Mark Feighn, is more complicated, but depends on the same tools, namely the index inequality of [GL95] and Theorem 5.2 of [Pau97]. The additional work needed is to consider the case that $\text{rank}(A_1) = \dots = \text{rank}(A_L)$ and to show, by a Stallings fold argument, that the first Betti number of the quotient of A_1 modulo the normal closure of A_i is a bounded, strictly increasing function of $i = 1, \dots, L$.

3.2 Geometric models and vertex group systems

Here is our main construction of vertex group systems.

Proposition 3.3. *Given a rotationless $\phi \in \text{Out}(F_n)$ represented by a CT $f: G \rightarrow G$, and a geometric stratum $H_r \subset G$, let X be a geometric model for H_r with accompanying notation as in Definition 2.4, in particular the surface S and the map $S \xrightarrow{j} Y \subset X$. Let $L \subset X$ be the complementary subgraph as in Definition 2.6. For any subcomplex $K \subset L$ such that $j(\partial S) \subset K$, the subgroup system $[\pi_1 K]$ is a vertex group system.*

Proof. We must construct an F_n -tree T with trivial arc stabilizers in which $[\pi_1 K]$ is realized. The idea is to use the Morgan-Shalen construction of dual trees to measured geodesic laminations [MS91], applied to a hyperbolic structure on the surface S . The strategy follows the diagram below: we must first pry S and L apart (the left square of the diagram), apply the Morgan-Shalen construction (the middle square), and then put the pieces back together

(the right square).

$$\begin{array}{ccccccc}
& \widetilde{X} & \xleftarrow{\tilde{q}} & \mathcal{P} & \xrightarrow{r} & \mathcal{F} & \longrightarrow T \\
\downarrow & & \downarrow & & \downarrow & & \downarrow \\
X & \xleftarrow{q} & L \amalg S & \longrightarrow & \left(\mathcal{F}/F_n \right) & \longrightarrow & \left(T/F_n \right)
\end{array}$$

The left column — the universal covering of X — is known. On the top row all objects will be F_n -actions and all arrows will be F_n -equivariant maps. All vertical arrows will be quotients modulo F_n . The **R**-tree T , and the **R**-forest \mathcal{F} which precedes it in the order of construction, will be acted on indiscretely by F_n , and their quotients are shown below them, with the parentheses to emphasize their non-Hausdorff nature. It remains to describe, in order, q , \mathcal{P} , \mathcal{F} , and T and various arrows.

The “prying apart” step is accomplished by the quotient map $q: L \amalg S \rightarrow X$ that is defined as follows: $q|L$ is just the inclusion $L \subset X$; and $q|S = j|S$. The identifications made by q are: the circles $\partial_0 S \subset S$ and $j(\partial_0 S) \subset L$ are identified; for $1 \leq i \leq m$ each point $x \in \partial_i S$ is identified with $\gamma_i(x) \in G_{r-1} \subset L$; and for each attaching point $x \in j(\text{int}(S)) \cap L$, the copies of x in L and in $\text{int}(S) \subset S$ are identified.

The space \mathcal{P} is simply a pushout, namely, the set of ordered pairs in the Cartesian product $\widetilde{X} \times (L \amalg S)$ whose two coordinates have the same image in X . Projections to the two Cartesian factors define the leftward and downward arrows out of \mathcal{P} . The action $F_n \curvearrowright \widetilde{X}$ induces an action $F_n \curvearrowright \mathcal{P}$ and the arrow $\tilde{q}: \mathcal{P} \rightarrow \widetilde{X}$ is F_n -equivariant. Letting L_1, \dots, L_A be the components of L , the arrow $\mathcal{P} \mapsto L \amalg S$ restricts on each component of \mathcal{P} to a universal covering map over a certain component L_1, \dots, L_A, S of $L \amalg S$. To be precise, fix an appropriate choice of base points and paths in X in order to identify $\pi_1(S), \pi_1(L_1), \dots, \pi_1(L_A)$ with subgroups of $\pi_1(X) \approx F_n$. Letting $\mathcal{L}(S)$ be a choice of left coset representatives of the subgroup $\pi_1(S) < F_n$, the components of \mathcal{P} over S are indexed as \widetilde{S}_c for $c \in \mathcal{L}(S)$, such that $\text{Stab}(\widetilde{S}_c) = \pi_1(S)^c = c^{-1} \pi_1(S) c < F_n$. Similarly, the components of \mathcal{P} over each L_a are indexed as $\widetilde{L}_{a,c}$, for c in a set $\mathcal{L}(L_a)$ of left coset representatives of $\pi_1(L_a) < F_n$, so that $\text{Stab}(\widetilde{L}_{a,c}) = \pi_1(L_a)^c$ (some components L_a may be contractible in which case $\mathcal{L}(L_a) = F_n$).

In what follows we shall abuse notation by identifying the components $\widetilde{S}_c, \widetilde{L}_{a,c}$ of \mathcal{P} with their embedded images $\tilde{q}(\widetilde{S}_c), \tilde{q}(\widetilde{L}_{a,c})$ in \widetilde{X} .

The space \mathcal{F} will be an F_n -forest obtained as a quotient of $F_n \curvearrowright \mathcal{P}$ by collapsing each component $\widetilde{S}_c \subset \mathcal{P}$ to an **R**-tree, using the dual **R**-tree construction [MS91], and by collapsing each component $\widetilde{L}_{a,c} \subset \mathcal{P}$ to a simplicial **R**-tree with appropriate vertex stabilizers.

Fix a hyperbolic structure with totally geodesic boundary on S (below we shall alter this structure by isotopy in order to establish the important property (*)). Fix a geodesic lamination $\Lambda \subset S$ which fills S and a transverse measure μ with support Λ . For example, one could take Λ, μ to be the unstable measured geodesic lamination of the pseudo-Anosov mapping class on S that is part of the definition of the geodesic model. The outcome of the construction in [MS91] is a $\pi_1(S)$ -tree $T(\Lambda)$ dual to Λ , as follows. Lift Λ to a geodesic lamination $\widetilde{\Lambda} \subset \widetilde{S}$ on the universal covering space \widetilde{S} of S . Lift μ to a transverse measure $\tilde{\mu}$

with support $\tilde{\Lambda}$. There is a continuous pseudo-metric $\tilde{d}(x, y) = \text{int}_{[x, y]} \mu$ where $[x, y]$ is the geodesic segment connecting $x, y \in \tilde{S}$ (to construct \tilde{d} we use, in part, that Λ fills, for if Λ had a compact leaf that that leaf would be an atom of μ and this definition of \tilde{d} would fail to produce a continuous pseudo-metric). Recalling that a principal region of $\tilde{\Lambda}$ in \tilde{S} is the closure of a component of $\tilde{S} - \tilde{\Lambda}$, the value of $\tilde{d}(x, y)$ depends only on the principal regions or nonboundary leaves of $\tilde{\Lambda}$ containing x, y respectively, and furthermore $\tilde{d}(x, y) = 0$ if and only if x, y are contained in the same principal region or nonboundary leaf. Let $\tilde{S} \rightarrow T(\Lambda)$ be the metric quotient of the pseudo-metric \tilde{d} , collapsing to a point each principal region and each nonboundary leaf of $\tilde{\Lambda}$, inducing a metric d on $T(\Lambda)$ such that the quotient map $\tilde{S} \rightarrow T(\Lambda)$ is an isometry from \tilde{d} to d . These objects in \tilde{S} — namely, $\tilde{\Lambda}, \tilde{\mu}, \tilde{d}$ — are all equivariant with respect to the deck transformation action $\pi_1(S) \curvearrowright \tilde{S}$, and therefore there is an induced isometric action $\pi_1(S) \curvearrowright T(\Lambda)$. As proved in [MS91], the space $T(\Lambda)$ with the metric d is an **R** tree, the action $\pi_1(S) \curvearrowright T(\Lambda)$ is minimal and has trivial arc stabilizers. Furthermore, the nontrivial point stabilizers can be described in any of the following ways: they are the subgroups of $\pi_1(S)$ conjugate to the fundamental groups of crown principal regions of Λ in S ; they are the infinite cyclic groups conjugate to the fundamental groups of the boundary circles $\partial_0 S, \dots, \partial_m S$; and they are the stabilizers of the components of $\partial \tilde{S}$.

Define a continuous pseudometric to each component of \mathcal{P} as follows. On each component \tilde{S}_c over S we use the pseudo-metric described above, arising from the lift of the measured foliation (Λ, μ) on S . Also, for each component L_a , $a = 1, \dots, A$, we choose a path pseudo-metric which assigns length 0 to each edge of $L_a \cap K$ and length 1 to each edge of $L_a \setminus K$, and then on each component $\tilde{L}_{a,c}$ of \mathcal{P} over $L_{a,c}$ we lift the path pseudometric of L_a . These pseudo-metrics are equivariant with respect to the deck transformation action $F_n \curvearrowright \mathcal{P}$, in other words the action restricts to pseudometric isometries amongst the components of \mathcal{P} .

Define $r: \mathcal{P} \rightarrow \mathcal{F}$ to be the F_n -equivariant quotient which on each component \tilde{S}_c or $\tilde{L}_{a,c}$ of \mathcal{P} restricts to the metric quotient of the pseudo-metric on that component, producing a component of \mathcal{F} which is an **R**-tree denoted $r(\tilde{S}_c)$ or $r(\tilde{L}_{a,c})$, respectively. The action $F_n \curvearrowright \mathcal{P}$ descends via r to an action $F_n \curvearrowright \mathcal{F}$. The metrics on the components of \mathcal{F} are equivariant with respect to this action, i.e. the action restricts to metric isometries amongst the components of \mathcal{F} . By the Morgan-Shalen theorem, each action $\pi_1(S)^c \curvearrowright r(\tilde{S}_c)$ is a minimal **R**-tree action whose vertex stabilizers are the conjugates in F_n of the subgroups $\pi_1(\partial_0 S), \dots, \pi_1(\partial_m S)$ that are contained in $\pi_1(S)^c$. Also, each action $\pi_1(L_a)^c \curvearrowright r(\tilde{L}_{a,c})$ is a Bass-Serre tree for the graph of groups structure on the quotient of L_a obtained by collapsing to a point each component of $L_a \cap K$, and the vertex stabilizers of this action are the subgroups of F_n contained in $\pi_1(L_a)^c$ that are conjugate to $\pi_1(K_j)$, for each each component K_j of K that is contained in L_a . Putting this altogether, it follows that the nontrivial point stabilizers of the **R**-forest action $F_n \curvearrowright \mathcal{F}$ are precisely the subgroups conjugate to the infinite cyclic groups $\pi_1(\partial_0 S), \dots, \pi_1(\partial_m S)$ and the subgroups in the subgroup system $[\pi_1 K]$.

We shall refer to the components of \mathcal{F} of the form $r(\tilde{S}_c)$ as the *S-components*, and to their points as *S-points* of \mathcal{F} . Also, we refer to components of the form $r(\tilde{L}_{a,c})$ as *L-components* and their points as *L-points* of \mathcal{F} . In this language, under the action $F_n \curvearrowright \mathcal{F}$, the subgroup system of stabilizers of all *L*-points of \mathcal{F} is exactly $[\pi_1 K]$, and the subgroup system of stabilizers of all *S* points of \mathcal{F} is exactly $\{[\pi_1 \partial_0 S], \dots, [\pi_1 \partial_m S]\}$.

The F_n -tree T will be constructed from \mathcal{F} by identifying certain pairs of points in an F_n -equivariant manner. But we must not identify two L -points of \mathcal{F} with nontrivial stabilizers for that will create a point of T whose stabilizer is larger than any subgroup in the system $[\pi_1 K]$. We shall avoid identifying L -points altogether by imposing an additional constraint on our construction.

We may assume, after possibly isotoping the hyperbolic structure on S by an isotopy supported on a compact subset of $\text{int}(S)$, that the following holds:

Property (*) No attaching point lies on the j -image of a crown principal region of Λ in S .

No two attaching points lie on the j -image of same nonboundary leaf of Λ or the same ideal polygon principal region in S .

This is possible because Λ has infinitely many leaves but only finitely many boundary leaves in S . The finite subset of attaching points may then be put into bijective correspondence with some finite subset of $j(\text{int}(S))$ that does satisfy (*), and there is an isotopy of S as desired taking one of these subsets to the other.

The tree T is the quotient of \mathcal{F} modulo the F_n -invariant partition generated by the following F_n -invariant relation: given $\sigma, \lambda \in \mathcal{F}$, an S -point and an L -point respectively, we say that σ, λ are *incident* if the subset $q(r^{-1}(\sigma)) \cap q(r^{-1}(\lambda)) \subset \tilde{X}$ is nonempty. Observe that as a consequence of Property (*), every S -point is incident to at most one L -point; the quotient map $\mathcal{F} \mapsto T$ is therefore injective on the set of L -points. Furthermore, from Property (*) it follows that if σ, λ as above are incident then one of the following two cases holds. In the first case, $q(r^{-1}(\sigma)) \cap q(r^{-1}(\lambda)) \subset \tilde{X}$ is a single point covering an attaching point of X , which occurs only if there exist $c \in \mathcal{L}(S)$ and $c' \in \mathcal{L}(L_a)$, for some $a = 1, \dots, A$, such that $r^{-1}(\sigma) \subset \tilde{S}_c$ is either an ideal polygon principal region or nonboundary leaf, and $r^{-1}(\lambda) \subset \tilde{L}_{a,c'}$ is either a universal covering of a component of $L_a \cap K$ or a single point lying over a point of $L_a - (L_a \cap K)$; furthermore, in this first case the stabilizer of σ is trivial. In the second case, $q(r^{-1}(\sigma)) \cap q(r^{-1}(\lambda))$ is a component of $\partial \tilde{S}_c$ for some $c \in \mathcal{L}(S)$, which occurs only if $r^{-1}(\sigma)$ is a universal covering over some crown principal region of S_c and $r^{-1}(\lambda) \subset \tilde{L}_{a,c'}$ is a universal covering of some component of $L_a \cap K$; furthermore, in this second case the stabilizer of σ is in the subgroup system $\{[\pi_1 \partial_0 S], \dots, [\pi_1 \partial_m S]\}$, and it is a subgroup of the stabilizer of λ which is in the subgroup system $[\pi_1 K]$. Note also that because of the inclusion $j(\partial S) \subset K$, every S -point σ with nontrivial stabilizer is incident to some L -point λ and the second case holds. Putting together this description of incidence and the stabilizers involved, it follows that the stabilizers of points of T are exactly the subgroups in the subgroup system $[\pi_1 K]$, as required.

Note that it is indeed possible for an L -point $\lambda \in \mathcal{F}$ to be incident to more than one S -point: this happens when $r^{-1}(\lambda) \subset \tilde{L}_{a,c}$ is a universal covering of some component of $L_a \cap K$ which contains at least two attaching points, or the j -images of at least two components of $\partial_i S$, or at least one of each; this, however, is inconsequential to our present purposes.

So far, T has been described as a set on which F_n acts, with point stabilizers being the required subgroup system $[\pi_1 K]$. Furthermore, the point stabilized by a particular subgroup in the system $[\pi_1 K]$ is unique. It follows that *any* \mathbf{R} -tree action $F_n \curvearrowright T$ has trivial edge stabilizers. It remains to define an \mathbf{R} -tree metric on T for which the action $F_n \curvearrowright T$ is isometric and minimal.

The metric on T is defined as follows. There is a partially defined metric d' on T which

is the union of the pushforwards via the map $\mathcal{F} \mapsto T$ of the metrics on the components of \mathcal{F} . Define $d(x, y) = \inf \left(\sum_{i=0}^I d'(x_{i-1}, x_i) \right)$ where the infimum is taken over all *connecting sequences* from x to y , meaning a sequence of the form $x = x_0, x_1, \dots, x_I = y$ such that for each $i = 1, \dots, I$ the pair x_{i-1}, x_i is contained in the image of some component of \mathcal{F} . The definition of d is evidently F_n -equivariant and so F_n acts by isometries of d . It is evident, at least, that d is a pseudo-metric.

To check that T is an **R**-tree we need to see that the components of \mathcal{F} are connected up in a tree-like pattern under the map $\mathcal{F} \mapsto T$, in the following sense. Recalling $H_r \subset G \subset X$ in the data of the geometric model X , let $F_n \curvearrowright U$ be the Bass-Serre tree of the graph of groups obtained from G by collapsing to a point each component of $K \cup H_r$, and note that this graph of groups may also be obtained from X by collapsing to a point each component of the union of the total lifts of K and S . There is an induced map $s: T \mapsto U$ whose restriction to each **R**-tree $r(\tilde{S}_c)$ is a constant map onto the point $s(r(\tilde{S}_c))$ which is the image under $\tilde{X} \mapsto U$ of $\tilde{S}_c \subset \tilde{X}$, and whose restriction to each subtree $r(\tilde{L}_{a,c})$ is a $\pi_1(L_a)^c$ -equivariant embedding onto the subtree $s(r(\tilde{L}_{a,c}))$ of U which equals the image under the map $\tilde{X} \rightarrow U$ of $\tilde{L}_{a,c} \subset \tilde{X}$. The map $s: T \mapsto U$ is, set theoretically, simply the quotient that collapses each subtree of the form $r(\tilde{S}_c)$ to a point $s(r(\tilde{S}_c))$. The map s therefore has the following effect on the subtrees of T of the form $r(\tilde{S}_c)$ and $r(\tilde{L}_{a,c})$. First, two subtrees $r(\tilde{S}_c), r(\tilde{L}_{a,c'}) \subset T$ have nontrivial intersection in T if and only if their s -images in U have nontrivial intersection if and only if $s(r(\tilde{S}_c))$ is a point of the subtree $s(r(\tilde{L}_{a,c'}))$. Second, two subtrees $r(\tilde{L}_{a,c}), r(\tilde{L}_{a',c'})$ have s -images with nontrivial intersection in U if and only if they intersect a common subtree $r(\tilde{S}_{c''}) \subset T$ if and only if $s(r(\tilde{S}_{c''})) \in s(r(\tilde{L}_{a,c})) \cap s(r(\tilde{L}_{a',c'}))$, in which case the latter intersection consists solely of the point $s(r(\tilde{S}_{c''}))$.

We now define the arc $[x, y] \subset T$ with any endpoints $x, y \in T$ and verify that T is an **R**-tree. Assuming $sx \neq sy$, the arc $[sx, xy] \subset U$ is a unique concatenation of subarcs $[sx, sy] = [\xi_0, \xi_1] * \dots * [\xi_{J-1}, \xi_J]$, $J \geq 1$ such that each $[\xi_{j-1}, \xi_j]$ is a maximal subarc contained in some subtree $s(r(\tilde{L}_{k_j, c_j}))$, $j = 1, \dots, J$. Each such subarc pulls back uniquely to a subarc $[x_{j-1}, x'_j] \subset r(\tilde{L}_{k_j, c_j})$. Letting $x'_0 = x$ and $x_J = y$, for each $j = 0, \dots, J$ we have two points x'_j, x_j in the same subtree $r(\tilde{S}_{c'_j})$, connected by a unique arc $[x'_j, x_j] \subset r(\tilde{S}_{c'_j})$. We therefore obtain an embedded arc in T which is expressed as a concatenation without backtracking of the form

$$[x, y] = [x'_0, x_0] * [x_0, x'_1] * \dots * [x_{J-1}, x'_J] * [x'_J, x_J]$$

whose first and last terms may be degenerate but whose other terms are all nondegenerate. This expression for $[x, y]$ makes sense even when $sx = sy$ but $x \neq y$, in which case $J = 0$ and $x = x'_0, y = x_0 \in r(\tilde{S}_{c'_0})$. From the form of the metric on T we have

$$d(x, y) \leq d'(x'_0, x_0) + d'(x_0, x'_1) + \dots + d'(x_{J-1}, x'_J) + d'(x'_J, x_J)$$

whose first and last terms may be zero but whose other terms are all nonzero. But in fact, this inequality is an equation, because *any* connecting sequence from x to y must contain $x = x'_0, x_0, x'_1, \dots, x_{J-1}, x'_J, x_J$ as an increasing subsequence. The arc $[x, y]$ is therefore isometrically embedded in T . And it is the unique embedded arc in T connecting x to y ,

because any such arc must contain $x = x'_0, x_0, x'_1, \dots, x_{J-1}, x'_J, x_J$ as an increasing subsequence, and the arcs between successive points are unique in the corresponding subtrees. This shows that T is an \mathbf{R} -tree.

Finally, minimality of the action $F_n \curvearrowright T$ is an immediate consequence of minimality of the actions $F_n \curvearrowright U$ and $\pi_1(S)^c \curvearrowright r(\bar{S}_c)$ for each $c \in \mathcal{L}(\pi_1 S)$.

This completes the proof of Proposition 3.3. \square

References

- [BFH97] M. Bestvina, M. Feighn, and M. Handel, *Laminations, trees, and irreducible automorphisms of free groups*, Geom. Funct. Anal. **7** (1997), 215–244.
- [BFH00] ———, *The Tits alternative for $\text{Out}(F_n)$. I. Dynamics of exponentially-growing automorphisms.*, Ann. of Math. **151** (2000), no. 2, 517–623.
- [BFH04] ———, *Solvable subgroups of $\text{Out}(F_n)$ are virtually Abelian*, Geometriae Dedicata **104** (2004), 71–96.
- [BH92] M. Bestvina and M. Handel, *Train tracks and automorphisms of free groups*, Ann. of Math. **135** (1992), 1–51.
- [CB88] A. Casson and S. Bleiler, *Automorphisms of surfaces after Nielsen and Thurston*, London Mathematical Society Student Texts, vol. 9, Cambridge University Press, 1988.
- [Coo87] D. Cooper, *Automorphisms of free groups have finitely generated fixed point sets*, J. Algebra **111** (1987), no. 2, 453–456.
- [FH11] M. Feighn and H. Handel, *The recognition theorem for $\text{Out}(F_n)$* , Groups Geom. Dyn. **5** (2011), 39–106, Preprint.
- [FLP⁺79] A. Fathi, F. Laudenbach, V. Poenaru, et al., *Travaux de Thurston sur les surfaces*, Astérisque, vol. 66–67, Société Mathématique de France, 1979.
- [FM12] B. Farb and D. Margalit, *A primer on mapping class groups*, Princeton Mathematical Series, vol. 49, Princeton University Press, 2012.
- [GJLL98] D. Gabai, A. Jaeger, G. Levitt, and M. Lustig, *An index for counting fixed points of automorphisms of free groups*, Duke Math. J. **93** (1998), no. 3, 425–452.
- [GL95] D. Gabai and G. Levitt, *The rank of actions on \mathbf{R} -trees*, Ann. Sci. École Norm. Sup. (4) **28** (1995), no. 5, 549–570.
- [Haw08] T. Hawkins, *Continued fractions and the origins of the Perron-Frobenius theorem*, Arch. Hist. Exact Sci. **62** (2008), no. 6, 655–717. MR MR2457065 (2009i:01018)
- [HM09] M. Handel and L. Mosher, *Subgroup classification in $\text{Out}(F_n)$* , arXiv:0908:1255, 2009.

- [HM13a] ———, *Subgroup decomposition in $\text{Out}(F_n)$: Introduction and Research Announcement*, arXiv:1302.2681, 2013.
- [HM13b] ———, *Subgroup decomposition in $\text{Out}(F_n)$, Part II: A relative Kolchin theorem*, arXiv:1302.2379, 2013.
- [HM13c] ———, *Subgroup decomposition in $\text{Out}(F_n)$, Part III: Weak Attraction Theory*, in preparation, 2013.
- [HM13d] ———, *Subgroup decomposition in $\text{Out}(F_n)$, Part IV: Relatively irreducible subgroups*, in preparation, 2013.
- [LL08] G. Levitt and M. Lustig, *Automorphisms of free groups have asymptotically periodic dynamics*, J. Reine Angew. Math. **619** (2008), 1–36. MR MR2414945
- [Mil82] R. T. Miller, *Geodesic laminations from Nielsen’s viewpoint*, Adv. in Math. **45** (1982), no. 2, 189–212.
- [MS91] J. Morgan and P. Shalen, *Free actions of surface groups on \mathbf{R} -trees*, Topology **30** (1991), no. 2, 143–154.
- [Pau97] F. Paulin, *Sur les automorphismes extérieurs des groupes hyperboliques*, Ann. Sci. École Norm. Sup. (4) **30** (1997), no. 2, 147–167. MR MR1432052 (98c:20070)
- [Ser80] J. P. Serre, *Trees*, Springer, New York, 1980.
- [Spa81] E. Spanier, *Algebraic topology*, Springer, New York—Berlin, 1981.
- [Sta00] J. Stallings, *Surfaces mapping to wedges of spaces. A topological variant of the Grushko-Neumann theorem*, Groups—Korea ’98 (Pusan), de Gruyter, 2000, pp. 345–353.
- [SW79] P. Scott and C. T. C. Wall, *Topological methods in group theory*, Homological group theory, Proceedings of Durham symposium, Sept. 1977, London Math. Soc. Lecture Notes, vol. 36, 1979, pp. 137–203.